Running head: GENERALIZATION AT RETRIEVAL

Generalization at retrieval using associative networks with transient weight changes

Kevin D. Shabahang (k.shabahang@gmail.com)¹

Hyungwook Yim (hwyim@hanyang.ac.kr)^{1,2}

Simon J. Dennis (simon.dennis@unimelb.edu.au)¹

¹School of Psychological Sciences, The University of Melbourne, Australia

²Department of Cognitive Sciences, Hanyang University, Seoul, Republic of Korea

Acknowledgements

We would like to thank Jeremy Caplan for his valuable feedback and comments on the manuscript.

Abstract

Without having seen a bigram like “her buffalo” before, you can easily tell that it is grammatical because “buffalo” can be aligned with more common nouns like “cat” or “dog” that have been seen in contexts like “her cat” or “her dog” -- the novel bigram structurally aligns with representations in memory. We present a new class of associative nets we call *Dynamic-Eigen-Nets*, and provide simulations that show how they generalize to patterns that are structurally aligned with the training domain. Linear-Associative-Nets respond with the same pattern regardless of input, motivating the introduction of *saturation* to facilitate other response states. However, models using saturation cannot readily generalize to novel, but structurally aligned patterns. Dynamic-Eigen-Nets address this problem by dynamically biasing the eigenspectrum towards external input using temporary weight changes. We demonstrate how a two-slot Dynamic-Eigen-Net trained on a text corpus provides an account of bigram *judgement-of-grammaticality* and *lexical decision* tasks. We end with a simulation showing how a Dynamic-Eigen-Net is sensitive to syntactic violations introduced in bigrams, even after the association encoding those bigrams is deleted from memory. We propose Dynamic-Eigen-Nets as associative nets that generalize at retrieval, instead of encoding, through recurrent feedback.

Keywords: Pattern-Completion; Content Addressable Memory; Auto-associative; Recurrent; Short-Term-Plasticity; Generalization

1. Introduction

The ability to learn the structure of various grammars underlying serially ordered events, be it parsing a sentence, following a melody, or tying one's shoes, is a hallmark of intelligent behaviour. A learning theoretic account that can handle the problem of grammar learning must specify how statistical regularities derived from a small subset of grammatical strings can encode sufficient constraints for the system to recognize the entire set. Restricting ourselves to the linguistic domain¹, how can we tell what serial ordering of words is congruent and what serial ordering is incongruent without having seen all congruent permutations of words? When recognition is not an option, structural generalization is required to determine which words can follow which other words. How is it that learners can generalize the structure of an infinite domain based on a finite number of exemplars?

One approach is to abstract away the contextual details of the contents of experience during encoding to form a generic representation. Assuming that memories are stored as connectivity weights in a network of simple processing units, a *generalization-at-encoding* view (e.g., Hinton, 1990) implies that at the time of storage (and perhaps also during sleep, Stickgold & Walker, 2013) gradual changes to the connectivity of the network fine-tune the system for transforming its input into some desired output. The transformation compresses (“summarizes”) the input into a lower-dimensional representation that retains information relevant to the mapping. Some difference between the generated pattern and a desired pattern is then used to obtain an error. The error is used to adjust the connectivity weights in the direction that reduces the error, and the whole cycle is repeated until the error is minimized. If the input was every word in a sentence, except for one target word that was treated as the desired output, over many iterations with many sentences, words of similar syntactic and semantic classes would cluster together in the compressed space (e.g., Westbury & Hollis, 2019). If the bigram “you know” was never seen, but “I know” and “they know” were seen, the system can align the representations and infer that “you know” is congruent since “you” is very similar to “they” and “I” and “know” and both “I” and “know” can precede “know”.

The mapping of input into a latent space brings out a higher-order co-occurrence structure that is not explicit in the original space (Landauer & Dumais, 1997; Grefenstette, 1994). In *word2vec* (Mikolov et al., 2013), words surrounding a target word are compressed into a low-dimensional representation that

¹ In principle, processing any serially ordered domain may be modelled using the mechanisms we propose, however, restricting the domain to linguistic utterances makes the system's dynamics easier to follow.

is used to predict the target word itself. The network is optimized to make the mapping with small subsequences (*contexts*) sampled from a large text-corpus. Words that co-occur in the same contexts (e.g., “squirrel” and “chipmunk”) map to similar representations in the lower-dimensional space.

Error-based neural architectures are widely employed in cognitive models to formalize various theories of human information processing (see Rogers & McClelland, 2014, for a review), but they have some limitations. Compressing the input into a lower-dimensional representation renders encoded patterns more vulnerable to cross-talk. When a sufficient number of new patterns are encoded, interference can completely wipe out previously encoded representations (McCloskey & Cohen, 1989; Mannering & Jones, 2020). Interference can be reduced by making the changes to the weights very small for a given learning trial (i.e., each forward pass of an input through the network followed by a backward propagation of error). When the weight changes are small, the network requires about an order of magnitude more learning trials (around 600) to match human learners (around 10 trials) on the same task (McCloskey & Cohen, 1989), and interference is still far greater in the network than what is observed with human learners (Barnes & Underwood, 1959)².

An alternative *generalization-at-retrieval* view (Hintzman, 1986; see Jones, 2018, for a recent overview) is that instead of assuming a mapping of the input into some latent space, a noisy copy is directly encoded into the connectivity structure. The presentation of a cue activates some of the processing-units, which then initiate a cascade of signals that reverberate through the network until the mutual constraints imposed by the connectivity structure dynamically resolve to complete retrieval (Hintzman, 1986). The encoded patterns are not abstractions. Abstraction takes place as retrieval stabilizes to a pattern formed through the integration of all previously encoded patterns in resonance with the retrieval cue.

While generalization at retrieval is not mutually exclusive with generalization during encoding, it has been argued that deferring generalization until retrieval better meets the flexibility demands of the environment (Hintzman, 1988). One reason is that the system has access to the relevant cues when attempting to generalize, hence escaping the need to prospectively assume a generic form that can meet later requirements. During retrieval, constraints (i.e. memories) that are relevant to the immediate circumstances can be selected and integrated online. The absence of dimensionality reduction allows the

² Another way the problem is mitigated in the *word2vec* algorithm is by randomly sampling contexts from the training corpus instead of sequentially processing it from start-to-finish. Making the assumption for human learning would be akin to chopping up experience into mutually incoherent fragments.

system to freely add new constraints with little immediate concern for how they will impact performance on other representations. In instance-based models like MINERVA 2 (Hintzman, 1984, 1986), each new instance of experience is stored by appending a noisy copy to the end of a rectangular matrix, say E ; the width of the matrix grows with experience. At retrieval, a cue is compared with every instance in memory, in parallel, and a retrieved pattern is constructed by summing all instances into a single vector, each weighted by its similarity to the retrieval cue.

Jamieson and Mewhort (2009, 2010, 2011; Chubala & Jamieson, 2013) demonstrated that a MINERVA 2 model is capable of distinguishing grammatical and ungrammatical strings in artificial grammar tasks, where learners are presented with a set of strings generated from a probabilistic Finite State Automaton, and are later queried to discriminate between strings that are generated from the same rules or violate the rules in some way. Dennis (2005; also see Kwantes, 2005) provided a large-scale instance-based model, trained on natural language text, that captures various phenomena in semantic composition. Johns and Jones (2015) have provided simulations that capture various phenomena in the sentence reading literature by equipping an instance-based model with semantic vectors constructed using the BEAGLE algorithm (Jones & Mewhort, 2007; for more recent results see Johns, Jamieson, Crump, Jones & Mewhort, 2020).

An alternative class of models, called *associative nets*³, employ a generalization-at-retrieval approach without adding memory patterns to a limitless memory store. By having away with a compressed representation (i.e., a hidden layer), associative nets are also distinct from error-driven neural networks. Associative nets provide a mechanistic account of memory retrieval through the collective action of a mass of associations. They have been used to model phenomena as varied as categorical perception (Anderson et al., 1977), serial recall (Farrell & Lewandowsky, 2002), reading comprehension (Kintsch, 1998), and letter perception (McClelland & Rumelhart, 1981).

Following Hebb (1949), associative nets assume that synapses between pairs of “neurons” strengthen when the neurons activate within a short time interval. If a single neuron encodes a single feature in a stimulus (e.g., colour or size), then the strengthened synapses encode co-occurrence patterns between features that compose the stimulus. Stimulus patterns are represented as vectors, where the value in each element stands for the rate of firing of a particular neuron⁴. A geometric interpretation of

³ Associative nets do not neatly fall into the instance-based versus prototype distinction in the previous work.

⁴ In some contexts it may be more accurate to think of each unit in the network as a bundle of neurons, as is done in mean-field approximations, however, it is more straightforward to take each to be a single unit without loss of generality.

the state vector is to view it as a point in a high-dimensional space, making them dynamic analogues to semantic spaces (e.g., BEAGLE). In a classic Hebbian associative net, instead of appending representations of experience to a wide matrix, as in MINERVA 2, patterns are encoded in proportion to the sum of their outer products, each superimposed into a weight matrix of fixed dimensionality. In essence, the Hebb rule binds coactive features into an assembly of neurons that can mutually excite one-another later during retrieval. If we treat the rectangular memory used in MINERVA 2 as the collection of experiences to which the system is exposed, memory in an associative net is proportional to the product of the experience matrix with its transpose, $\mathbf{W} \propto \mathbf{E}\mathbf{E}^T$ (cf the co-occurrence matrix).

Associative nets have suffered from one of two problems. On the one hand, systems like the Interactive Activation and Competition model (Rumelhart & McClelland, 1986) and Construction Integration networks (Kintsch, 1998) require handcrafted connection weights, and have been difficult to scale beyond toy problems meant to demonstrate their key properties. On the other hand, associative networks like the Hopfield net (Hopfield, 1982) and the Brain-State-in-a-Box (Anderson, Silverstein, Ritz & Jones, 1977) have been restricted to memory retrieval tasks; they do not generalize beyond the reconstruction of encoded patterns. In this paper, we propose a mechanism that extends the generalization capabilities of associative networks to allow them to more effectively compete with instance-based models. More specifically, by generalization we mean *structural generalization*: deciding if permutations of symbols come from a combinatorial domain, based on experience with just a small subset of instances. We use the terms “structural generalization” and “combinatorial generalization” interchangeably throughout the manuscript, and will often opt to simply use “generalization”. We demonstrate how generalization during retrieval may accomplish a simplified grammar induction task: judging the serial order congruity of word pairs.

The network dynamics in an associative net are defined by a first-order difference equation -- a recurrence relation. The recurrence drives retrieval by specifying how the weight matrix, $\mathbf{W}$, and the momentary state vector, $\mathbf{x}_t^T$, interact to yield the next momentary state vector, $\mathbf{x}_{t+1}^T$. Geometrically, the recurrence relation defines a law of motion in high-dimensional feature-space. During retrieval, an input probe, $\mathbf{x}_0^T$ initializes the state for the first time-point. The state at the next time-point is a function, f , of the vector-matrix multiplication of the current state and the weight matrix, $\mathbf{x}_{t+1}^T = f(\mathbf{x}_t^T \mathbf{W})$. In the Linear-Associative-Net, the function f is simply unit-normalization. The process is carried out iteratively until further cycles have no additional effect on the state vector (i.e., when $\mathbf{x}_{t+1}^T \approx \mathbf{x}_t^T$). Recurrence

forces the co-occurrence between features encoded from previously learned patterns to interact with the probe until the system reaches an equilibrium (steady) state -- the retrieved pattern.

The weight matrix can be decomposed into a set of orthogonal dimensions that capture the degree of change corresponding to different aspects of experience. Multiplying a vector with a matrix and unit-normalizing, as done with the Linear-Associative-Net, pushes the vector closer to the direction of variance that captures the maximum amount of variance in the encoded experience -- the dominant eigenvector of $\mathbf{W}$. Because a Linear-Associative-Net always settles to the dominant eigenvector, it is only capable of generating a single response. In order to increase the size of the response set, Anderson et al. (1977; Anderson, 1995) introduced saturation in the Brain-State-in-a-Box by bounding each neuron's activation by a constant. Saturation constrains possible states of activation within a box and forces convergence to one of the corners. The bounding box halts the system before the state gets dominated by the lead eigenvector, therefore allowing a larger number of steady states -- responses.

The Brain-State-in-a-Box assumes the encoded patterns are corners of a hypercube (i.e., mutually orthogonal bi-polar vectors with an equal number of 1's and -1's). When the encoded patterns are orthogonal, they become the eigenvectors of the weight matrix. Therefore each encoded memory forms its own eigenvector in the Brain-State-in-a-Box. Although saturation enables a larger possible set of responses, the Brain-State-in-a-Box is limited because it restricts steady states to single eigenvectors corresponding to previously encoded patterns. A desirable property for a system capable of combinatorial generalization is the ability to combine constraints from multiple eigenvectors, but this is not possible in the Brain State in a Box.

Previous attempts have been made to remedy the generalizability of associative nets, but they have been limited when dealing with correlated structure among memory patterns. Strategies used to overcome those limitations depend on further complicating assumptions. Amari (1977) explored the effect of encoding correlated patterns in associative nets, where non-linearity was introduced by using a binary-threshold activation function. The output of each element of the state-vector was set to one, if, and only if, it passed some fixed threshold, and was zero otherwise. Using several variations of the Hebb rule, Amari showed that fine-tuning the threshold parameter can mitigate against noise introduced by the cross-talk between correlated patterns, however, the cross-talk quickly became unwieldy when additional noise was introduced at encoding. The noise intolerance suggests the system will not scale to deal with an input stream resembling raw experience.

Here, we assume that input to the system enhances the conductance between connections that link the input elements to each other, in addition to affecting the initial state. For instance, if you are shown words such as “the cat sat on the mat”, the connections linking all the words in the sentence to each other are assumed to be facilitated for the duration of the time the pair is maintained in working memory. The facilitation forces the emergence of a neural assembly that forces a positive feedback loop between the input elements, a reverberatory loop. The increase in conductance linking the input elements forces the system to reach equilibrium to a point in state-space that is near the input. As the system approaches equilibrium, it is pulled towards some of the eigenvectors in the matrix and pushed away from others. Equilibrium is therefore a compromise between the input and the structure encoded along the eigenvectors. In the model we present, we assume that each time the system is probed, a neural assembly is temporarily superimposed over the static weights that were learned during previous encoding. The input-driven connectivity changes are not permanent; they remain present while the system is approaching equilibrium and are reset after.

We explore the system in relation to the eigenvectors of the weight matrix and show how transient assemblies re-weight various eigenvectors (directions of variance) encoded in memory, dampening the impact of some and magnifying the impact of others. We show that transient assemblies provide a linear alternative to saturation for overcoming the dominant eigenvector problem. Whereas in classic associative nets the weights combining the eigenvectors -- the eigenvalues -- remain static, a transient assembly dynamically weights the eigenvectors based on each eigenvector’s similarity to the input pattern. Therefore we call the system a Dynamic-Eigen-Net. The additional gain-control on the eigenstructure prevents saturation to a single eigenvector and enables mixed-eigenstates, or retrieved states that are spread across multiple eigenvectors.

We first provide some toy simulations to demonstrate the essential characteristics of Dynamic-Eigen-Nets. We then scale up the system to encode bigram information from a naturalistic text corpus. We use the scaled-up Dynamic-Eigen-Net to make contact with some data from lexical decision tasks that manipulate syntactic congruity between adjacent word-present pairs and show priming between syntactically congruent pairs. Noting the paucity of experiments that use bigrams to manipulate syntactic variables, we collected our own data in a 2AFC bigram acceptability task. As we show later, our model’s performance provides a good match to human data. We end by testing the system’s generalization ability, where we delete associations between word-pairs (e.g., association between “her”

to “buffalo”) and check if the network can still distinguish between the grammatical form (“her buffalo”) and its ungrammatical counterpart (“she buffalo”).

2. Properties of associative nets

We now illustrate some key properties of three variants of associative networks. The first variant is a simple Linear-Associative-Net, the second is the Brain-State-in-a-Box, and the third is a Dynamic-Eigen-Net.

The encoded patterns and retrieval cues are identical across the three associative net variants; we only change the algorithm driving the system towards equilibrium following a probe -- the recurrence relation. The Linear-Associative-Net and the Dynamic-Eigen-Net do not make any particular assumptions about representation, but the Brain-State-in-a-Box requires that memories correspond to mutually orthogonal bi-polar vectors that are corners of a hyper-cube. We define four orthogonal vectors, each of dimensionality four, and assign each a single word in English. We use $[1, 1, 1, 1]^T$ for “the”, $[-1, 1, -1, 1]^T$ for “cat”, $[1, -1, -1, 1]^T$ for “a”, and $[1, 1, -1, -1]^T$ for “dog”. The capital letter “T” superscript stands for the transposition operation (i.e., swapping rows for columns or vice-versa). We assume that input to the system is an eight-dimensional vector, and construct bigram vectors by concatenating pairs of individual word vectors. Hence, bigrams are encoded with the word in the first serial position active in the first slot and the word in the second serial position active in the second slot.

We construct the bigrams “the cat” and “a dog” by concatenating the respective word-vectors, in sequence. We encode the two bigram vectors with unequal strengths to show how the Linear-Associative-Net always responds with the stronger pattern, even when the other pattern is only fractionally weaker. First we unit-normalize each bigram vector and then encode it into memory. We set the strength of the stronger pattern to 1.2 and the strength of the weaker pattern to 1.1. For all following simulations in the toy demonstrations, we assume a weight matrix,

$$\mathbf{W} = 1.2\mathbf{m}_{the-cat}\mathbf{m}_{the-cat}^T + 1.1\mathbf{m}_{a-dog}\mathbf{m}_{a-dog}^T$$

where $\mathbf{m}_{the-cat}$ and $\mathbf{m}_{a-dog}$ are the bigram vectors corresponding to “the cat” and “a dog”, respectively.

To initialize a probe for retrieval, we take an 8-dimensional vector with one, or more, slots populated with individual word vectors, add noise to it, and unit-normalize. We add noise because

noise-tolerance is a necessary property of biological neural networks. Gaussian noise is superimposed onto the pattern via, ϵ , an 8-dimensional vector of samples from a zero-mean Gaussian distribution with its standard deviation, $\sigma = 0.3$. We set the probe as,

$$\mathbf{x}_0 = \frac{\mathbf{x} + \epsilon}{\|\mathbf{x} + \epsilon\|} \quad \text{where, } \epsilon = \mathcal{N}(0, \sigma)$$

The vector, $\mathbf{x}$, has dimensionality 8. The double vertical bars denote vector length.

We characterize the system's state at each time-point in terms of the level of activation for the primitive word vectors, yielding four activation values in the first slot and four activation values in the second slot. To compute the activation values for each slot, we segment the state vector from the middle and take the first half as the first slot and the second half as the second slot. We set the activation value for a word in position one (or two) as the absolute value (*cf* Farrell et al., 2002) of the vector cosine of its primitive vector and the first slot (or second).

2.1 Linear Associative Net

In a Linear-Associative-Net, we define the recurrence relation in terms of the unit-normalized vector-matrix multiplication of the current state vector and the weight matrix. It is given by,

$$\mathbf{x}_{t+1}^T = \frac{\mathbf{x}_t^T \mathbf{W}}{\|\mathbf{x}_t^T \mathbf{W}\|}$$

We designate a small criterion, ϵ , and terminate retrieval when the change from one time-point to the next falls below criterion, i.e., $\|\mathbf{x}_{t+1}^T - \mathbf{x}_t^T\| < \epsilon$. We set the convergence criterion to 1e-07 for all the following toy simulations.

Rewriting the weight matrix, $\mathbf{W}$, in terms of its eigenstructure is a common way to analyze the encoded structure. The weight matrix can be decomposed into a superposition of the outer-products of its eigenvectors, each weighted by their respective eigenvalues, λ_i , $\mathbf{W} = \Sigma(\lambda_i \hat{\mathbf{e}}_i \hat{\mathbf{e}}_i^T)$, where $\hat{\mathbf{e}}_i$ stands for the i 'th eigenvector of $\mathbf{W}$. Dropping the normalization, we then have,

$$\mathbf{x}_{t+1}^T = \mathbf{x}_t^T \mathbf{W} = \mathbf{x}_{t-1}^T \mathbf{W} \mathbf{W} = \mathbf{x}_0^T \mathbf{W}^t = \sum_{i=1}^N ((\lambda_i^t \mathbf{x}_0^T \hat{\mathbf{e}}_i) \hat{\mathbf{e}}_i^T)$$

We include parentheses around the term, $\lambda_i^t \mathbf{x}_0^T \hat{\mathbf{e}}_i$, to emphasize that it is a scalar. Because of the exponent over the eigenvalues, and because the eigenvalues are monotonically decreasing in magnitude, from the first to the last, in the limit, as we multiply the vector with the matrix and unit-normalize, $\mathbf{x}_{t+1}^T$

converges to the dominant eigenvector, $\hat{\mathbf{e}}_{\max}$, except in the case when $\mathbf{x}_0^T \hat{\mathbf{e}}_{\max} = 0$.⁵ If we assume any level of noise, the Linear-Associative-Net is a single-response memory system because $\mathbf{x}_0^T \hat{\mathbf{e}}_{\max}$ will rarely equal zero.

Table 1 shows the probability of each response (along the columns) as a function of each probe (along the rows). The system always converges to the dominant pattern (“the cat”) regardless of the input. We provide the probabilities based on 1000 runs in Tables 1, 2, and 3. Each run is different because of the noise term, ϵ , added to the probe.

Table 1

Probabilities of response across probes show how a Linear-Associative-Net is restricted to a single response

	Probability of response					
	the cat	a dog	the dog	a cat	dog the	cat a
Probe						
the cat	1	0	0	0	0	0
a dog	1	0	0	0	0	0
the _	1	0	0	0	0	0
a _	1	0	0	0	0	0
the dog	1	0	0	0	0	0
a cat	1	0	0	0	0	0
dog the	1	0	0	0	0	0
cat a	1	0	0	0	0	0

note. the underscores preceded by determiners, “the” and “a”, stand for an empty second slot (i.e., all elements populated with Gaussian noise).

2.2 Brain-State-in-a-Box

⁵ The process of iteratively multiplying a vector with a matrix and unit-normalizing is referred to as the Power Iteration and is the basis for computing eigenvectors and eigenvalues numerically.

The Brain-State-in-a-Box remedies the dominant eigenvector problem by introducing saturation, forcing a maximum and minimum over the activation range. The modified recurrence relation is given by,

$$\mathbf{x}_{t+1}^T = S(\mathbf{x}_t^T \mathbf{W})$$

where

$$S(\mathbf{x}_i) = \begin{cases} 1 & \mathbf{x}_i \geq 1 \\ \mathbf{x}_i & -1 < \mathbf{x}_i < 1 \\ -1 & \mathbf{x}_i \leq -1 \end{cases}$$

with 1 being the saturation constant. Normalizing is no longer used with saturation.

Table 2 has the same form as Table 1, and shows the probability of responses (columns) as a function of different probes (rows). Whereas the Linear-Associative-Net always settled to the dominant pattern, “the cat”, the Brain-State-in-a-Box settles to “the cat” when probed with “the cat” and “the _”, and settles to “a dog” when probed with “a dog” or “a _”. Because of the noise and the increased basin of attraction -- gravitational pull around the point corresponding to the encoded pattern -- for “the cat”, probing the system with the weaker pattern settles to “the cat” about half the time. The partial probes demonstrate the pattern-completion capabilities of the system, again showing a preference for the stronger pattern when probed with the partial pattern aligning with the weaker bigram. The underscore denotes an empty slot; the elements are populated by Gaussian noise.

The second two probes, “the dog” and “a cat”, are novel combinations of the primitive word vectors. Both “the” and “a” have been encoded in the first slot and both “dog” and “cat” have been encoded in the second slot. Therefore they structurally align with the encoded patterns. Despite their alignment with the eigenstructure of the system, the responses generated by the Brain-State-in-a-Box are restricted to previously stored patterns, with a strong preference for the stronger pattern, shown in the fifth and sixth rows. The final two probes, “dog the” and “cat a”, are the same novel combinations, except that the two words have been swapped so that they are no longer aligned with the encoded structure. The system always responds with “the cat” with either probe. Critically, since the system settles to a previously encoded pattern regardless of the structural alignment of novel probes with previously encoded patterns, the Brain-State-in-a-Box cannot structurally generalize.

Table 2

Probabilities of response across probes show how a Brain-State-in-a-Box is restricted to previously stored patterns

	Probability of response					
	the cat	a dog	the dog	a cat	dog the	cat a
Probe						
the cat	1	0	0	0	0	0
a dog	.551	.449	0	0	0	0
the _	1	0	0	0	0	0
a _	.507	.493	0	0	0	0
the dog	.998	.002	0	0	0	0
a cat	.996	.004	0	0	0	0
dog the	1	0	0	0	0	0
cat a	1	0	0	0	0	0

note. the underscores preceded by determiners, “the” and “a”, stand for an empty second slot (i.e., all elements are populated by Gaussian noise).

2.3 Dynamic-Eigen-Net

We now show how transient assemblies enable the system to generalize to novel patterns based on combinations of multiple eigenvectors. We model transient assemblies by assuming a temporary increase in the weights corresponding to the active entries in the input. We assume that the temporary change follows the presentation of the input and persists for the duration of the subsequent set of iterations. After convergence, the weights are reset. We model transient assemblies by superimposing the outer-product of the probe, with itself, into the weight matrix. We also normalize the state vector after each iteration, as was done with the Linear-Associative-Net. By getting rid of saturation, we are left with a linear system we can characterize directly in relation to the eigenstructure.

The recurrence relation for a Dynamic-Eigen-Net is given by,

$$\mathbf{x}_{t+1}^T = \frac{\mathbf{x}_t^T (\mathbf{W} + \mathbf{x}_0 \mathbf{x}_0^T)}{\|\mathbf{x}_t^T (\mathbf{W} + \mathbf{x}_0 \mathbf{x}_0^T)\|}$$

In both a Linear-Associative-Net and the Brain-State-in-a-Box, convergence filters out any component orthogonal to the encoded eigenvectors. In a Dynamic-Eigen-Net the input pattern persists and is integrated into other components that do align with the eigenstructure.

If we let $\mathbf{x}_\infty^T$ be the state in the limit, and λ_∞ be the primary eigenvalue of $\mathbf{W} + \mathbf{x}_0 \mathbf{x}_0^T$, the following equation describes the steady-state:

$$\mathbf{x}_\infty^T (\mathbf{W} + \mathbf{x}_0 \mathbf{x}_0^T) = \sum_{i=1}^N ((\lambda_i \mathbf{x}_\infty^T \hat{\mathbf{e}}_i) \hat{\mathbf{e}}_i^T) + (\mathbf{x}_\infty^T \mathbf{x}_0) \mathbf{x}_0^T = \lambda_\infty \mathbf{x}_\infty^T$$

The term, $\lambda_\infty \mathbf{x}_\infty^T$, follows from the fundamental eigenvalue identity because $\mathbf{x}_\infty^T$ is the primary eigenvector of $\mathbf{W} + \mathbf{x}_0 \mathbf{x}_0^T$. It should be emphasized that for the term, $\sum_i ((\lambda_i \mathbf{x}_\infty^T \hat{\mathbf{e}}_i) \hat{\mathbf{e}}_i^T)$, the eigenvectors and eigenvalues correspond to the original weight matrix before the outer-product of the initial pattern was added.

The parentheses around $\lambda_i \mathbf{x}_\infty^T \hat{\mathbf{e}}_i$ and $\mathbf{x}_\infty^T \mathbf{x}_0$ are added to emphasize that they are scalar terms. The activation pattern converges toward the direction of each of the eigenvectors, weighted by its dot-product with the current state and the eigenvalue, plus the initial pattern, weighted by its dot-product to the current state. Since the states are assumed to have unit-normal length, the dot-products correspond to vector cosines. In a nutshell, adding the outer-product into the weight matrix not only dynamically weights each eigenvector, but also forces the persistence of the initial pattern.

If we post-multiply both sides with $\mathbf{x}_\infty^T$, we obtain an equation for the dominant eigenvalue of the converged state, λ_∞ , which we will later use as a measure of familiarity:

$$\lambda_\infty = \sum_{i=1}^N ((\lambda_i \mathbf{x}_\infty^T \hat{\mathbf{e}}_i)^2) + (\mathbf{x}_\infty^T \mathbf{x}_0)^2$$

In the limit, the lead eigenvalue is the squared sum of the similarity of the state and each eigenvector, weighted by the corresponding eigenvalue, plus the square of the similarity of the state and the initial probe⁶.

In some implementations of the Brain-State-in-a-Box (e.g., Anderson, 1995) the initial state is also included in the recurrence relation. We can include the initial state, as in:

⁶ The lead eigenvalue is similar to Smolensky's (1986) measure of harmony and the global strength used in MINERVA2 (Hintzman, 1986; 1988).

$$\mathbf{x}_{t+1}^T = \mathbf{x}_t^T \mathbf{W} + \mathbf{x}_0^T$$

But the rule is only equivalent to our update rule for the first iteration, i.e.,

$$\mathbf{x}_1^T = \mathbf{x}_0^T (\mathbf{W} + \mathbf{x}_0 \mathbf{x}_0^T) = \mathbf{x}_0^T \mathbf{W} + \mathbf{x}_0^T$$

In the second iteration, for the Dynamic-Eigen-Net, we have,

$$\mathbf{x}_2^T = (\mathbf{x}_0^T \mathbf{W} + \mathbf{x}_0^T) (\mathbf{W} + \mathbf{x}_0 \mathbf{x}_0^T) = \sum_{i=1}^N ((\lambda_i^2 + \lambda_i) \mathbf{x}_0^T \hat{\mathbf{e}}_i \hat{\mathbf{e}}_i^T) + \left(\sum_{i=1}^N (\lambda_i (\mathbf{x}_0^T \hat{\mathbf{e}}_i)^2) + 1 \right) \mathbf{x}_0^T$$

If we were to simply add the initial state to the output, we would have,

$$\mathbf{x}_2^T = (\mathbf{x}_0^T \mathbf{W} + \mathbf{x}_0^T) \mathbf{W} + \mathbf{x}_0^T = \sum_{i=1}^N ((\lambda_i^2 + \lambda_i) \mathbf{x}_0^T \hat{\mathbf{e}}_i \hat{\mathbf{e}}_i^T) + \mathbf{x}_0^T$$

For the second iteration, the Dynamic-Eigen-Net includes an additional weight for the initial state, i.e., $\sum_i (\lambda_i (\mathbf{x}_0^T \hat{\mathbf{e}}_i)^2)$, which enables the similarity between the initial state and the eigenspectrum of the weight matrix to adjust the relative weight given to the initial state. In general, each additional iteration in the Dynamic-Eigen-Net introduces a new high-order mixture of weights that correspond to interactions between the initial state and the eigenstructure. It is the high-order mixture terms that drive the compromise between the eigenstructure and the input during recurrence.

Table 3 shows the probability of responses as a function of the probe for a Dynamic-Eigen-Net. The first two probes (“the cat” and “a dog”) yield a similar pattern of response as the Brain-State-in-a-Box. One notable difference is that the Dynamic-Eigen-Net is not limited to previously stored patterns (e.g., sometimes it responds with “the dog” when probed with “a dog”). As in the Brain-State-in-a-Box, the Dynamic-Eigen-Net is sensitive to the relative strength of the encoded items. The system settles to “a dog”, the weaker pattern, with a less than one probability whereas when probed with the stronger pattern, it always settles to “the cat”. In contrast to the Brain-State-in-a-Box, the Dynamic-Eigen-Net settles to the weaker pattern with much greater probability. Sensitivity to differences in the encoded strengths gives the system a way to track probabilities in the environment because in a continuously learning system, higher probability signals will have an increased basin of attraction.

The pattern completion dynamics are evident in the next two probes, “the _” and “a _”. The system always completes the stronger pattern in accordance with the encoded bigram (i.e., “the _” to “the cat”), but it sometimes completes the partial probe corresponding to the weaker pattern (“a _”) with the noun that was in the stronger pattern (“a cat”) and sometimes overrides the partial probe with the

stronger pattern entirely (i.e., “a _” settles to “the cat”). Whereas the Brain-State-in-a-Box almost always favoured the strong pattern, the Dynamic-Eigen-Net settles to the strong pattern with a much lower probability. Critically, it sometimes even settles to a completely novel pattern (“a cat”).

Table 3

Probabilities of response across probes show how a
Dynamic-Eigen-Net accommodates novel patterns

	Probability of response					
	the cat	a dog	the dog	a cat	dog the	cat a
Probe						
the cat	1	0	0	0	0	0
a dog	.006	.862	.132	0	0	0
the _	1	0	0	0	0	0
a _	.003	.991	0	.006	0	0
the dog	.006	.024	.97	0	0	0
a cat	.058	.009	0	.933	0	0
dog the	1	0	0	0	0	0
cat a	1	0	0	0	0	0

note. the underscores preceded by determiners, “the” and “a”, stand for an empty second slot (i.e., all populated by Gaussian noise).

The probes “the dog” and “a cat” in Table 3 show the response probabilities for the structurally aligned novel bigrams. Although the system sometimes converges to the originally encoded patterns, it settles to the novel inputs with high probability, hence generalizing to structurally consistent patterns. Finally, the last two probes, “dog the” and “cat a”, show the probabilities of responding when the words in the novel bigrams swap serial-positions, resulting in novel combinations that misalign with the encoded structure. The system never settles to the misaligned patterns (the columns corresponding to “dog the” and “cat a” have zero probability). In the case of misaligned novel patterns, the system defaults to the strongest memory: “the cat”. Hence, although the system accommodates novel patterns

that structurally align with the encoded patterns, it suppresses misaligned patterns through interaction terms between the initial state and the eigenstructure.

To illustrate the dynamics, Figure 1 shows the activations of the words as a function of feedback iteration, for two example runs corresponding to the novel probes, “a cat” (panel A) and “cat a” (panel B). The pattern of responses shown in Figure 1 corresponds to cases with the highest probability responses in Table 3. The plots on the left side correspond to the first slot whereas the ones on the right correspond to the second slot. When the novel pattern aligns with previously encoded structure (panel A), the system settles to the novel pattern, however, when the same words in the bigram are swapped to break the alignment (panel B), the system defaults to the strongest memory item that does align (i.e., “the cat”). We have included the familiarity values corresponding to each cue (λ_{ω}). If we were to use a criterion for recognition based on familiarity, we may say that the system accepts strings that are part of the learned language, but rejects those that are not. Because of its ability to distinguish structurally aligned novel patterns from those that are misaligned, the Dynamic-Eigen-Net surpasses the generative complexity of the Brain-State-in-a-Box, while still respecting the constraints encoded in memory.

Figure 1. Probing a Dynamic-Eigen-Net with a novel bigram yields a response corresponding to the initial input, only if the probe is aligned with the structure encoded in the system (panel A) but not when it is misaligned (panel B). Note also that the global familiarity, λ_{∞} , is larger when the pattern aligns relative to when it does not align.

3. From Theory to Data

Having illustrated that the Dynamic-Eigen-Net is capable of generalizing in a toy demonstration, we scale up the model by training it on a text corpus meant to approximate the variable and unstructured input experienced by human observers throughout the first eighteen years of development. To preserve the simplicity of representation, we continue with a two-slot model as before. The subsequent simulations are not meant as complete cognitive models, but are presented to showcase the model’s generalization capability at scale. A more complete model will require further architectural assumptions that are beyond the scope of the present manuscript.

We consider two tasks that can be easily modified to deal with bigrams: the lexical decision task and the judgement-of-grammaticality task. The lexical decision task requires participants to decide whether strings of letters with which they are presented are words or nonwords. The judgement-of-grammaticality task requires subjects to decide whether a sequence of words forms a well-formed utterance or not.

The major question we address in the following simulations is whether encoding pairwise serial-order associations into a Dynamic-Eigen-Net yields differences in familiarity between grammatical and ungrammatical bigrams. If a serial-order association exists in memory between the pair of words in a grammatical bigram and not for the corresponding ungrammatical bigram, then a difference in familiarity simply demonstrates recognition. To show that the system can generalize, we must demonstrate that it is able to exploit the structure in the input to yield higher familiarity for grammatical strings over ungrammatical strings, without any knowledge of the specific grammatical strings queried. In our final simulation, we delete the associations, both forward and backward⁷, between the words in the syntactically congruent bigram before comparing its familiarity to the corresponding incongruent bigram.

3.1. *Scaling up the Dynamic-Eigen-Net*

We made several modifications to the Dynamic-Eigen-Net presented in the toy examples in order to scale the system. The first modification was a change from distributed-codes to local-codes. We used distributed-codes for the toy demos to meet the Brain-State-in-a-Box’s representation assumptions.

⁷ We remind the readers here that the weight matrix is symmetric, such that, $\mathbf{W}_{ij} = \mathbf{W}_{ji}$.

Instead of representing words as vectors of -1's and 1's, we represented them with vectors with all, but one, element set to zero. The nonzero element, at an index unique to each word, was set to one. Local-codes ensure orthogonality and facilitate an increase in the vocabulary size as they enable the use of sparse matrix data structures. If we encode a set of two-slot local-code vectors from a corpus using Hebbian learning, we obtain a co-occurrence matrix containing the number of times pairs of words were adjacent in the corpus. With the exception of one set of simulations, we obtain the counts by sliding a two-word window across the TASA corpus to encode all the lag-one sequential dependencies, similar to the toy examples. To simulate one of the experiments, we used a French corpus. Specifically, we used a subset of a POS-tagged French Wikipedia corpus called WikipediaFR2008 (<http://redac.univ-tlse2.fr/corpus/wikipedia/wikipediaFR-2008-06-18.tag.7z>).

The raw co-occurrence matrix does not work directly as the memory system to the Dynamic-Eigen-Net, especially when the input domain follows a Zipfian distribution. That is, the most frequent word is about twice as frequent as the second-most frequent word, and so forth for the third-most frequent word relative to the second-most and so on. Suppose we know the joint probability of two words, A and B , and we want to know if the two words are associated. If the two words are independent, then their joint-probability will simply be the product of their marginal probabilities. Hence under the independence assumption, the ratio of their joint-probability and the product of their marginal probabilities will be near one. We can take the logarithm of the ratio so that when the two words are independent, their association value will be zero. If the two words have a higher-than-chance probability of occurring together, the ratio will be larger than one, and the logarithm of the ratio will be positive. In contrast, if the two words have a less-than-chance probability of co-occurring, the ratio will be less than one, in which case the logarithm of the ratio will be negative. The measure of association just described is called the pointwise mutual information (PMI; Church & Hanks, 1990).

The PMI is related to the BCM rule (Bienenstock et al., 1982), where synapses are assumed to be potentiated if the firing rate of a postsynaptic neuron surpasses a threshold, and attenuated otherwise. If we assume that the threshold at each synapse connecting pairs of neurons is proportional to the product of their base-rate firing, then the PMI has qualitative resemblance to the steady-state of weights learned using some variant of the BCM rule. The Dynamic-Eigen-Net does not require our specific choice of normalization. What is required is that the encoded associative structure captures the statistical

regularities of the input-space and that the eigenvalues of the corresponding weight matrix do not fall off very sharply from the largest eigenvalue to the smallest.

In a two slot concatenation model, the co-occurrence matrix can be partitioned into four submatrices. The top-left submatrix maps each word in the first slot to itself as does the lower-right submatrix for each word in the second slot. The top-right submatrix encodes the number of times a word in the second slot follows another word in the first slot, i.e., it yields the forward serial dependency counts. The bottom-left submatrix encodes the backward serial dependencies and is the transpose of the top-right submatrix. Let $\mathbf{C}_{ij}^{pq}$ be the co-occurrence count between the i 'th and j 'th word, corresponding to the (p, q) 'th submatrix. The submatrix indices for the top-left, top-right, bottom-left, and bottom-right are $(0, 0)$, $(0, 1)$, $(1, 0)$, and $(1, 1)$, respectively. The i and j indices range between zero and V , where V is the number of unique terms in the studied corpus. The full weight matrix, $\mathbf{W}$, for a two-slot model is therefore $2V$ by $2V$. Because the logarithm is not defined at zero, we add a small value, a , to all co-occurrence counts before applying PMI. We set the (i, j) 'th cell in the (p, q) 'th submatrix of the weights as,

$$\mathbf{W}_{ij}^{pq} = \log_2 \left(\frac{p_{ij}}{p_i p_j} \right)$$

where,

$$\begin{aligned} p_{ij} &= \frac{\mathbf{C}_{ij}^{pq} + \alpha}{T + \alpha V^2} \\ p_i &= \frac{\sum_{j=1}^V (\mathbf{C}_{ij}^{pq}) + \alpha V}{T + \alpha V^2} \\ p_j &= \frac{\sum_{i=1}^V (\mathbf{C}_{ij}^{pq}) + \alpha V}{T + \alpha V^2} \\ T &= \sum_{j=1}^V \sum_{i=1}^V \mathbf{C}_{ij}^{pq} \end{aligned}$$

We fixed a to 0.1 for the following simulations.

Figure 2 shows how the eigenvalue of the primary eigenvector towers over the other eigenvalues corresponding to the weight matrix after normalizing using the pointwise mutual information (PMI). The top twenty eigenvalues are shown.

One of the limitations of the PMI is that it does not sufficiently normalize against the Zipfian distribution of the input. Figure 2 shows the top twenty eigenvalues corresponding to the weight matrix, $\mathbf{W}$. The first eigenvalue is very large relative to the subsequent eigenvalues. In the Dynamic-Eigen-Net, we require that the transient weights are strong enough to exceed the dominant eigenvalue. We do so by dividing the entire matrix by the dominant eigenvalue and scaling the transient weights to 1.001, just slightly larger than one. As a consequence, the steady-state gets saturated by the dominant eigenvector and the transient weights, however, the variance captured by the lower order eigenvectors is necessary. The Dynamic-Eigen-Net is robust to differences in eigenvalue, but not when the eigenspectrum drops too sharply. We therefore reduce the strength of the dominant eigenvector by subtracting part of it from the weight matrix after we scale the matrix down by dividing by its largest eigenvalue,

$$\hat{\mathbf{W}} = \frac{1}{\lambda_{max}} \mathbf{W} - \eta \hat{\mathbf{e}}_{max} \hat{\mathbf{e}}_{max}^T$$

The corresponding recurrence equation used for the following simulations is,

$$\mathbf{x}_{t+1}^T = \frac{\mathbf{x}_t^T (\hat{\mathbf{W}} + \beta \mathbf{x}_0 \mathbf{x}_0^T)}{\|\mathbf{x}_t^T (\hat{\mathbf{W}} + \beta \mathbf{x}_0 \mathbf{x}_0^T)\|}$$

We set the dominant eigenvector inhibition parameter, η , to 0.55, and the scale for the transient weights, β , to 1.001. The convergence criterion was set to 1e-07. You can find the code for both the toy examples and the scaled-up simulations on our OSF page (<https://osf.io/g4axy/>).

Table 4 shows the steady-state activations for the top five most-active words in either the first- or second-slot, when the system probe was either in the second- or first-slot, respectively (*cf* Table 1 in Sahlgren et al., 2008 and Table 4 in Jones & Mewhort, 2007)⁸. By observation, the model does a good job of capturing serial-order dependencies.

Table 4

Cue	<cue> _	_ <cue>
king	's arthur midas minos george	luther dr. french rex the
president	nixon kennedy reagan johnson roosevelt	vice elected became former the
war	ii ended against broke i	civil world revolutionary vietnam indian
sea	level captain captains floor water	mediterranean above caribbean red black

Note. Each cue was either placed in the second, or first, slot and the words for the top five highest activations are shown. The underscore denotes an empty slot, which was filled with zeros.

In the following simulations, we compute a familiarity signal as the vector length of the state, prior to normalization, for the final recurrence iteration (i.e., λ_∞). We retain the familiarity signals across bigrams and subtract the familiarity for each invalid bigram from the familiarity for its paired valid bigram. Positive differences indicate greater familiarity for the valid bigram relative to the invalid bigram.

3.2. Goodman et al. (1997)

In one lexical decision task, Goodman, McClelland, and Gibbs (1981) found that participants were approximately 35 msc slower to decide whether a string of letters spells out a word, when it is syntactically inappropriate in the context of a preceding string of letters, that was also a word. Seidenberg et al. (1984) replicated the finding using the same stimuli. Goodman et al. (1981) provided a

⁸ One notable difference between our system and Sahlgren et al. (2008) and Jones and Mewhort (2007) is that we do not use a stop-list.

list of their stimuli, therefore we provide a direct simulation of their task. We matched their sample size and randomly constructed 32 different sets of valid-invalid bigram pairs using the method they describe. Figure 3 shows the distribution of familiarity differences for the 32 sets of bigram pairs. The vertical line marks the zero-point. The mass of the distribution falls above zero, showing that the Dynamic-Eigen-Net attributes greater familiarity to valid bigrams relative to invalid bigrams. Independent of congruency, Goodman et al. (1981) also reported slower response-times when the target word (second word in the bigram) was a verb relative to when it was a noun. We obtain greater familiarity when the target word is a noun relative to when it is a verb, $t(167) = 2.66$, $p = .008$, for the congruent bigrams. In contrast, for the incongruent bigrams, we obtain greater familiarity when the target word is a verb relative when it is a noun, $t(163) = -4.42$, $p < .001$.

Figure 3 shows how the familiarity for syntactically congruent bigrams are generally larger than the familiarity for the syntactically incongruent counterparts, using the same stimuli as Goodman et al. (1981). The model was trained on the TASA corpus.

3.3. Colé et al. (1994)

In a similar task, Colé and Segui (1994) presented subjects with pairs of letter-strings and asked them to respond “yes”, only if both were valid words in French. Subjects were slower to verify syntactically invalid bigrams relative to valid bigrams, with a greater response-time cost for bigrams whose initial word was closed-class relative to bigrams whose first word was open-class. Open-class words (e.g., adjectives, nouns, verbs etc.) correspond to content words and closed-class words (e.g., determiners, pronouns, prepositions) correspond to function words. The former group allows for the addition of new members (e.g., new nouns or adjectives), but the latter group does not.

Colé and Segui (1994) used French nouns as target words in their lexical decision task, because French nouns have inherent gender: some words are considered masculine and others feminine. For example, the word “chat” (“cat”) is a noun that is both singular and masculine. When preceded by the singular and masculine possessive pronoun, “mon” (“my”), or the singular and masculine adjective “joli” (“pretty”), the bigrams “mon chat” and “joli chat” are syntactically congruent because of the noun’s agreement in both gender and number. When “chat” is preceded by the plural pronoun, “mes”, or the feminine pronoun, “ma”, the resulting bigram disagrees in either number or gender, respectively. When “chat” is preceded by the singular and masculine adjective, “joli”, the corresponding bigram is syntactically congruent, but when it is preceded by the plural, “jolis”, or the feminine, “jolie”, then the bigram mismatches in either number or gender, respectively. The bigrams containing different possessive pronouns were part of the closed-class condition whereas the bigrams containing adjectives were part of the open-class condition. Because syntactic violations resulted in greater response-time costs for closed-class bigrams compared to open-class bigrams, Colé and Segui speculated that the two classes may be stored separately, an idea previously entertained by Garrett (1978). Model-derived differences in familiarity between grammatical and ungrammatical bigrams presented in our simulations provide an existence proof that distinct performance signatures can be obtained between the two classes without assuming qualitative differences in encoding or representation.

We used a French corpus to simulate Colé et al. (1994), and used the method described in Colé et al. (1994) to construct sets of paired valid and invalid bigrams. For each valid bigram we constructed two invalid bigrams with the first word mismatching the second word in either gender or number. First we collected 440 target words and constructed a closed-class valid-bigram and an open-class valid

bigram. We then created two invalid bigrams for each valid bigram by mismatching the context word in either gender or number.

Figure 4 shows the mean familiarity difference between matched congruent and incongruent bigrams, separately for closed-class, and open-class context words and number and gender mismatches. As shown by the positive familiarity differences, the model is sensitive to both number- and gender-violations for both closed- and open-class context words. When the context word is closed-class, the familiarity difference is reliably greater than when the context word is open-class, for gender mismatches, $t(439) = 5.28, p < .001$, but not for number mismatches, $t(439) = 1.08, p = .28$. Overall, the Dynamic-Eigen-Net explains the greater response-time costs for closed-class context words observed by Colé et al. (1994) based on greater differentiation in familiarity for bigrams containing closed-class words over those containing open-class words, however, it is unclear why the differentiation is only statistically reliable for the gender mismatches.

Figure 4 shows how the familiarity for the syntactically congruent bigrams is greater than the familiarity for corresponding syntactically incorrect bigrams, and how the mean familiarity difference is larger when the context word is a closed-class, relative to when it is open-class, using stimuli generated by Colé et al. (1994). The difference between the closed- and open-class bigrams is only statistically reliable for the bigrams that disagree in gender, but not for number. The model was trained on a subsample from the French Wikipedia corpus. The

error-bars show one standard-error around the mean.

4. bigram acceptability task

The syntactic congruity effect has also been shown in judgements of grammaticality. Münte et al. (1993) constructed bigrams composed of a pronoun followed either by a noun or a verb, such as “my cat” and “you think”. They introduced violations by swapping possessive pronouns with personal pronouns, or vice versa, as in “me cat” and “your think”. They then presented each bigram to participants, and asked them to decide whether it was grammatical or not. They found a response-time advantage when subjects responded to syntactically valid bigrams, relative to invalid bigrams. Münte et al. (1993) used English stimuli, but did not provide a list of their materials. Therefore we provide a qualitative match to their results. In order to compensate for the lack of data available for tasks that manipulate syntactic congruity using bigrams, we supplemented the manipulations employed by Münte et al. (1993) with seven additional syntactically congruent and incongruent pairs. Instead of using a yes-no judgement task as in Münte et al. (1993), we used a 2AFC design, where grammatical/ungrammatical bigram pairs were presented simultaneously and participants had to decide which was easier to read.

4.1. Participants

We recruited 20 native English speakers, with normal or corrected to normal vision, from *Amazon’s Mechanical Turk*. The participants were told about the task and provided consent prior to taking part. The study was approved by the Melbourne Psychological Sciences Board of Ethics.

4.2. Materials

We used the Stanford Part-of-speech (POS) tagger (Toutanova, Klein, Manning, & Singer, 2003) to tag words in the TASA corpus with their syntactic class; we did this in context and not with the words in isolation. We then grouped bigrams based on their syntactic composition. For each of nine different syntactic compositions, we used the 80 most frequent bigrams as our grammatical set. More specifically, we used the top 80 bigrams of each composition group, after discarding bigrams whose isolated POS

tags predicted by another tagger (Honnibal et al., 2016) conflicted with the Stanford tagger’s prediction. We also discarded bigrams that contained a word whose POS varied depending on context (e.g., “bear”). We then constructed an ungrammatical bigram for each bigram in the grammatical set. A listing of the bigrams is provided in the appendix, along with their corpus frequency and response-time and accuracy measures collected from the study.

The first column in Table 5 shows the valid (left) and invalid (right) bigram compositions used, each with an example in the corresponding row, in the second column. For example, in a bigram consisting of a possessive pronoun and a noun (PRP\$-NN), we can construct an ungrammatical bigram by turning the possessive pronoun into a personal pronoun (i.e., PRP-NN). If the PRP\$-NN bigram is “her cat”, the ungrammatical, PRP-NN, bigram would be “she cat”. For a determiner-noun bigram (DT-NN), such as “the cat”, we can construct a corresponding ungrammatical bigram, “cat the”, by swapping the two words (NN-DT). The PRP\$-NN versus PRP-NN and PRP-VBP versus PRP\$-VBP comparisons correspond to Münte et al. (1993), but the rest are novel violations.

Table 5

bigram composition violations along with an example for the nine sets of bigrams used in the acceptability task

Violation	Example
DT-NN vs NN-DT	“the cat” to “cat the”
PRP-VBP vs PRP\$-VBP	“you see” to “your see”
JJ-NN vs NN-JJ	“small amount” to “amount small”
IN-VBG vs IN-VBP	“of thinking” to “of think”
PRP\$-NN vs PRP-NN	“her cat” to “she cat”
NNS-VBP vs NN-VBP	“trees grow” to “tree grow”
VB-RBR vs RBR-VB	“learn more” to “more learn”
NN-VBZ vs NN-VBP	“cell divides” to “cell divide”
NN-IN vs IN-NN	“group of” to “of group”

note. Abbreviations correspond to the Penn Treebank convention for parts-of-speech. NN: singular or mass noun, IN: preposition, VB: verb, RBR: comparative adverb, VBZ: third-person, singular, and present tense verb, NNS: plural noun, VBP: non-third-person, present tense-verb, VBG: gerund or present-tense verb, DT: determiner, JJ: adjective, PRP: personal-pronoun, PRP\$: possessive pronoun.

4.3. Procedure

The study was conducted through the web-browser, using jsPsych (De Leeuw, 2015). Participants judged all pairs of bigrams from each of the nine violation types. Before making the judgments, we instructed participants to “put the index finger of your left hand on the 'A'/'a' key and the index finger of your right hand on the 'L'/'l' key”. We informed the participants that they would see pairs of bigrams presented on different sides of the screen, and that their job was to determine “which pair you find easier to read”. We asked that they “respond as fast as possible, while making sure you are also accurate”.

For each pair, the “correct” and “incorrect” bigrams were simultaneously displayed on opposite sides of the center of the screen. The bigrams were not too distant from the center (within about 5 degrees of visual angle at reasonable distances from the screen). The bigrams remained on screen until

the participant made a key press. The next bigram pair was immediately displayed following a response. Whenever the participants responded faster than 200 mscs or slower than 7500 mscs, we notified them that they were responding “too quickly” or “too slowly”, respectively. The position of the “correct” bigram was randomized.

We first had participants complete eighteen practice trials, using bigram pairs that were not used in the main study phase. After practice, we reminded them about their finger placement on the keyboard. Upon a final key press, we presented them 720 bigram pairs (9 x 80). All the bigrams were shuffled at random, separately for each participant.

4.4. Results and simulations

We excluded four participants from the analyses because their accuracy fell below 80%, leaving us with 16 participants. To assess the accuracy of the model, we tallied the proportion of times each valid bigram had a larger familiarity value than its corresponding invalid bigram.

The first column in Table 6 shows the nine different bigram comparison types, followed by the proportion of times the Dynamic-Eigen-Net favoured the grammatical over the ungrammatical bigrams. The third column shows the accuracy obtained from the 16 participants in the study. It is clear from Table 6 that performance for the Dynamic-Eigen-Net is at ceiling, with human performance lagging slightly behind. Interestingly, the two bigram comparisons for which the Dynamic-Eigen-Net did not perfectly discriminate, VB-RBR-to-RBR-VB and NN-IN-to-IN-NN, were also the least accurate among the participants.

For each comparison, containing 80 pairs of valid and invalid bigram exemplars, we computed a discriminability score for the Dynamic-Eigen-Net as the ratio of the mean familiarity difference and its standard deviation. Because we subtract the familiarity for the invalid bigrams from the valid bigrams to compute the familiarity difference, positive discriminability indicates greater familiarity for the valid relative to the invalid bigrams. We compared the discriminabilities from the model with the median response time taken for participants to judge pairs from the set of bigrams, separately for each bigram type. Figure 5 shows the median response time (RT) as a function of discriminability. We obtain a large correlation coefficient indicating that over 70% of the variance in the median RTs is captured by discriminability derived from the Dynamic-Eigen-Net. Overall participants were slowest to judge the NN-VBZ-to-NN-VBP, NN-IN-to-IN-NN, VB-RBR-to-RBR-VB, and NNS-VBP-to-NN-VBP pairs.

They were also the least discriminable for the Dynamic-Eigen-Net. In contrast, participants were fastest when making judgements about the JJ-NN-to-NN-JJ, IN-VBG-to-IN-VBP, PRP-VBP-to-PRP\$-VBP, PRPS-NN-to-PRP-NN, and DT-NN-to-NN-DT pairs. Consistent with the data, they were also the most discriminable for the Dynamic-Eigen-Net.

Table 6

Comparison	Dynamic Eigen Net	Human
DT-NN vs NN-DT	1	0.99
PRP-VBP vs PRP\$-VBP	1	0.94
JJ-NN vs NN-JJ	1	0.98
IN-VBG vs IN-VBP	1	0.97
PRP\$-NN vs PRP-NN	1	0.99
NNS-VBP vs NN-VBP	1	0.95
VB-RBR vs RBR-VB	0.975	0.93
NN-VBZ vs NN-VBP	1	0.95
NN-IN vs IN-NN	0.9625	0.86

Note. The bigram comparison types are sorted in descending order of discriminability. Abbreviations correspond to the Penn Treebank convention for parts-of-speech. NN: singular or mass noun, IN: preposition, VB: verb, RBR: comparative adverb, VBZ: third-person, singular, and present tense verb, NNS: plural noun, VBP: non-third-person, present tense-verb, VBG: gerund or present-tense verb, DT: determiner, JJ: adjective, PRP: personal-pronoun, PRP\$: possessive pronoun.

Figure 5 shows how the Dynamic-Eigen-Net’s discriminability between the syntactically congruent and incongruent bigram pairs, for each of the nine violations, is highly predictive of the median response time from participants in the bigram acceptability study. The model was trained on the TASA corpus. Abbreviations correspond to the Penn Treebank convention for parts-of-speech. NN: singular or mass noun, IN: preposition, VB: verb, RBR: comparative adverb, VBZ: third-person, singular, and present tense verb, NNS: plural noun, VBP: non-third-person, present tense-verb, VBG: gerund or present-tense verb, DT: determiner, JJ: adjective, PRP: personal-pronoun, PRP\$: possessive pronoun.

The discriminabilities we derive from the model correspond to the paired bigrams, where the invalid and valid bigrams are roughly matched by frequency since the invalid counterparts to the valid bigrams often have one or more words in common. It should be straight-forward for people to judge which of two, grammatical and ungrammatical, bigrams is grammatical without needing the two to closely match in their type of composition (e.g., “your know” vs “his dog”). To better understand how the model performs, we computed the discriminabilities between each of the nine grammatical bigram sets with each of the other ungrammatical bigram sets, yielding 81 different discriminability scores. For

each discriminability, we computed the familiarity difference between each bigram in the grammatical set with each bigram in the ungrammatical set, yielding $80^2 = 6400$ familiarity differences.

Figure 6 shows the discriminability between bigrams in each grammatical set compared with bigrams in each of the ungrammatical sets. Abbreviations correspond to the Penn Treebank convention for parts-of-speech. NN: singular or mass noun, IN: preposition, VB: verb, RBR: comparative adverb, VBZ: third-person, singular, and present tense verb, NNS: plural noun, VBP: non-third-person, present tense-verb, VBG: gerund or present-tense verb, DT: determiner, JJ: adjective, PRP: personal-pronoun, PRP\$: possessive pronoun.

Figure 6 shows a heatmap of the 81 discriminability scores comparing each of the nine grammatical bigrams (along the vertical axis) with each of the nine ungrammatical bigrams (along the horizontal axis). The cells shaded in a darker colour indicate higher discriminabilities, and those shaded very lightly indicate lower discriminabilities. The red font marks the cases where discriminability was higher for the ungrammatical set relative to the grammatical set. With the exception of four cells,

VB-RBR-vs-NN-DT, VB-RBR-vs-IN-NN, NNS-VBP-vs-IN-NN, and NN-VBZ-vs-IN-NN, the rest of the comparisons yield positive discriminability.

Our use of the term ungrammatical is out of convenience and only true in the sense that the members in the set very rarely occur in natural language. In some cases they only appear odd in a bigram. Some examples of the VB-RBR-vs-NN-DT comparisons include the bigram pairs, “tell more” vs “child the”, “go more” vs “money the”, and “mature more” vs “government the”. For the VB-RBR-vs-IN-NN comparisons, some examples include, “promote more” vs “for money”, “suffer more” vs “of price”, and “retain more” vs “of understanding”. Some example bigram pairs for the NN-VBZ-vs-IN-NN included “birds sing” vs “from home”, “symptoms include” vs “for food”, and “atoms combine” vs “of quality”. Finally, some examples for the NN-VBZ-vs-IN-NN bigrams include, “food gives” vs “of life”, “brain sends” vs “in life”, and “theory suggests” vs “of section”. The examples make it clear that, although rare, some of the ungrammatical bigrams may be plausible in larger sentential contexts. For example, the examples corresponding to the NN-DT bigrams may be plausible in larger utterances such as in “give the child the toy”, “when given the money the man laughed”, and “they told the government the truth”. Consistent with the heatmap and results from the bigram readability task, the IN-NN bigrams tend to have a large number of cases that can be considered grammatical in addition to bigrams that are very clearly ungrammatical. Some of the clearly ungrammatical bigrams include, “of advantage”, “of bit”, “of plenty”, and “of lot”.

We tallied the proportion of familiarity differences that were larger for the grammatical bigrams, relative to the ungrammatical bigrams, to examine performance when a bigram is randomly selected from one of the nine grammatical sets and compared to another randomly chosen bigram from one of the ungrammatical sets. The model chooses the grammatical bigram about 93% of the time. Hence the low discriminability values in the four comparisons are not fatal for the model. Based on Figure 6, the model predicts that the bigram comparisons with negative, or low, discriminability would be challenging to judge for human participants.

The frequency of the bigram in the corpus and the frequencies of the constituents composing the bigram have an impact on the discriminabilities. For instance, the DT-NN bigrams are the most discriminable from the ungrammatical bigrams. Our explorations of the eigenvectors of the weight matrix suggest that items falling into the DT class load strongly on the top eigenvectors. In general the loadings on the top eigenvectors favour closed-class words, which also explains how we obtained

greater familiarity differences for closed-class context words relative to open-class context words when modeling Colé et al. (1994). Closed-class words are not qualitatively distinct from the rest of the words, however, they form fixed dimensions of variation within the representational space. A key question is whether frequency information alone can predict the response times in the bigram acceptability task.

Table 7 shows the mean frequency of grammatical (column labeled with G_{bg}) and ungrammatical (UG_{bg}) bigrams in the TASA corpus, along with the frequency of the bigram constituents (G or UG subscripted with the position of the word: 1 or 2). To examine how much frequency information may be driving people's performance, the Pearson correlation between each frequency measure and the median RT from the bigram acceptability task is shown for each measure. Since all measures have a relatively high correlation with median RT, we regressed median RT on the six frequency measures in a single model. We observed that the mean frequency of the grammatical bigrams is the only statistically significant predictor: participants were faster to choose the grammatical bigrams over the ungrammatical bigrams, when the frequency for the grammatical bigram was higher in the corpus. The p -values for each of the corresponding regression weights from a single multiple-regression model are shown below the correlation coefficients.

Given that the mean bigram frequency for the grammatical bigrams predicts the median RT, it is possible that we can predict people's response times by simply using the frequency. In a final analysis, we regressed the median RT on both discriminability and mean grammatical bigram frequency, in a single model. Results showed that higher discriminability values reliably predicted faster RT, $t(6) = -3.477$, $p = 0.013$, and that any variance accounted for in median RT by the mean grammatical bigram frequency was no longer statistically significant, $t(6) = -1.300$, $p = 0.241$, when discriminability is included.⁹ Therefore we conclude that discriminability is sensitive to information that goes beyond frequency information, and that such information is highly predictive of people's processing speed in the context of the syntactic congruity effect.

⁹ Discriminability turned out to be the significant predictor, regardless of whether we placed it into the regression model first or second.

Table 7

Mean corpus frequency of the grammatical and ungrammatical bigrams and their constituents and their correlation with discriminability

Violation	G _{bg}	G ₁	G ₂	UG _{bg}	UG ₁	UG ₂
DT-NN vs NN-DT	1025.89	712384.86	3589.82	19.91	3589.82	712384.86
PRP-VBP vs PRP\$-VBP	254.26	67578.24	5111.06	0.14	24047.55	5111.06
JJ-NN vs NN-JJ	160.78	6301.44	4220.19	0.46	4220.19	6301.44
IN-VBG vs IN-VBP	75.55	174587.76	1542.26	0.49	174587.76	8373.55
PRP\$-NN vs PRP-NN	242.46	34800.04	4029.05	1.2	60463.53	4029.05
NNS-VBP vs NN-VBP	18.99	3301.61	3613.59	0.78	4171.66	3613.59
VB-RBR vs RBR-VB	17.34	2453.82	22522.22	0.29	22522.22	2453.82
NN-VBZ vs NN-VBP	13.55	3921.09	961.79	1.66	3921.09	3602.3
NN-IN vs IN-NN	539.06	2880.16	297245.33	97.47	297245.33	2880.16
Pearson's <i>r</i>	-.57	-.58	.19	.09	0	-.49
<i>p</i> -value	.033	.180	.313	.620	.163	.538

note. Abbreviations correspond to the Penn Treebank convention for parts-of-speech. NN: singular or mass noun, IN: preposition, VB: verb, RBR: comparative adverb, VBZ: third-person, singular, and present tense verb, NNS: plural noun, VBP: non-third-person, present tense-verb, VBG: gerund or present-tense verb, DT: determiner, JJ: adjective, PRP: personal-pronoun, PRP\$: possessive pronoun. G_{bg} denotes the grammatical bigram and UG_{bg} denotes the ungrammatical bigram, while G₁ corresponds to the first constituent of the grammatical bigram and G₂ corresponds to the second constituent. The same notation is used for the ungrammatical set. The Pearson's *r* corresponds to the correlation between each frequency measure and the median RT. The *p*-values are based on regression weights in a multiple-regression model, regressing median RT on all six predictors.

Generalization

So far the simulations we have discussed have been limited to recognition. That is, all of the associations connecting the words in the grammatical bigrams were encoded into memory. The system cannot rely on a simple match to memory to show generalization; it needs to rely purely on the structural properties encoded in the system. Suppose you are shown a pair of bigrams, “her buffalo” and “she buffalo”. If an association between “her”, in the first slot, and “buffalo”, in the second slot, exists in

memory and an association between “she”, in the first slot, and “buffalo”, in the second slot, does not exist or is much weaker, you can correctly pick out the grammatical bigram by relying on recognition. Suppose that we snip the association between “her” and “buffalo”. Then you can no longer rely on a simple recognition process, but must generalize. In the next simulation, we modelled the bigram discriminability task as before, but deleted the association between the two words making up each grammatical bigram, prior to probing the system with either the grammatical bigram or its ungrammatical counterpart. Before moving on to each new pair of grammatical and ungrammatical bigrams, we reset the association to its pre-lesioned state.

Table 8 shows the nine different violation types in the first column, an example for each violation in the second column, the discriminability obtained when the association for the grammatical bigram is left intact in the third column (i.e., recognition), and the discriminability obtained when the association for the grammatical bigram is deleted (“lesioned”), in the fourth column. Figure 7 shows the distribution of familiarity differences across the nine bigram violations, for the lesioned case. The vertical line marks the zero-point. In cases where discriminability is good, the mass of the familiarity differences lies clearly to the right of the zero-point.

As evident in Table 8, the model obtains positive discriminability for all, but one, bigram comparison set. The discriminabilities are all lower when the association is removed relative to when it is intact, showing the contribution of recognition to discriminability. Discriminability in the lesioned case cannot rely on recognition -- it is exploiting the global structure in the associations. Figure 7 echoes the discriminabilities in Table 8, with the mass of the familiarity differences falling closer to the zero-point for diminishing values of discriminability.

Out of the nine comparisons, the system shows good generalization for the VB-RBR-to-RBR-VB, PRP\$-NN-to-PRP-NN, IN-VBG-to-IN-VBP, DT-NN-to-NN-DT, JJ-NN-to-NN-JJ, and the PRP-VBP-to-PRP\$-VBP violations. The mass of the corresponding familiarity differences are markedly shifted towards the right of the zero-point. The system cannot generalize to the NN-VBZ-to-NN-VBP violations and can only partially generalize to the NNS-VBP-to-NN-VBP and NN-IN-to-IN-NN violations. In relation to human performance, the poorly generalized violations also appear to be the violations that people find more difficult to judge. As we have mentioned earlier, they are also the comparisons where the bigrams in the ungrammatical set are sometimes ambiguous as to whether they can be used in utterances. In relation to the closed- and open-class distinction, with the

exception of the NN-IN-to-IN-NN violations, which show poor generalization, all other violations that include closed-class words show good discriminability even after lesioning. Overall, the system is able to exploit the structure encoded in the eigenvectors to generalize without relying on recognition.

Table 8

The model shows good discriminability between valid and invalid bigrams across various syntactic violations

Violation	Example	Intact	Lesioned
DT-NN vs NN-DT	“the cat” to “cat the”	6.5	4.29
PRP-VBP vs PRP\$-VBP	“you see” to “your see”	5.74	4.39
JJ-NN vs NN-JJ	“small amount” to “amount small”	5.42	1.67
IN-VBG vs IN-VBP	“of thinking” to “of think”	5.3	3.38
PRP\$-NN vs PRP-NN	“her cat” to “she cat”	5.12	2.69
NNS-VBP vs NN-VBP	“trees grow” to “tree grow”	2.75	0.46
VB-RBR vs RBR-VB	“learn more” to “more learn”	2.31	2.52
NN-VBZ vs NN-VBP	“cell divides” to “cell divide”	2.08	−0.23
NN-IN vs IN-NN	“group of” to “of group”	1.55	0.52

note. The bigrams are sorted in descending order of the intact discriminability in the intact condition.

Abbreviations correspond to the Penn Treebank convention for parts-of-speech. NN: singular or mass noun, IN: preposition, VB: verb, RBR: adverb (comparative), VBZ: verb, (third-person singular present tense), NNS: plural noun, VBP: non-third-person verb (singular present tense), VBG: gerund verb (present participle), DT: determiner, JJ: adjective, PRP: personal-pronoun, PRP\$: possessive pronoun.

Figure 8. Distributions of familiarity differences (grammatical minus the ungrammatical) across the nine bigram comparison types, after deleting the associations for the grammatical bigrams, show that the Dynamic-Eigen-Net can generalize well for most of the comparisons. The red vertical line marks the zero-point. The distributions that fall on the right of the zero-point indicate good generalizability. Abbreviations correspond to the Penn Treebank convention for parts-of-speech. NN: singular or mass noun, IN: preposition, VB: verb, RBR: adverb (comparative), VBZ: verb, (third-person singular present tense), NNS: plural noun, VBP: non-third-person verb (singular present tense), VBG: gerund verb (present participle), DT: determiner, JJ: adjective, PRP: personal-pronoun, PRP\$: possessive pronoun.

How then is the system generalizing? In the first iteration the term¹⁰ in the first slot activates terms in the second slot, and the term in the second slot activates terms in the first slot. On each following iteration, the terms that had become activated in a previous iteration either inhibit or activate other terms in the opposite slot. For instance, given the probe “you know”, the term “you”, in the first slot, may activate terms such as “re” and “think” in the second slot while the term “know”, in the second slot, may activate terms such as “nt”, “i”, and “we”, in the first slot. When the two terms form an incongruent pair, as in “your know”, the terms activated by the input terms (i.e., “your” and “know”) are more likely to inhibit one-another. For instance, the term “your” in the first slot may activate “body” and “own” in the second slot. Because “body” and “own” are inconsistent with words activated in the first slot by the word “know”, they are more likely to be mutually inhibitory than facilitatory.

Figure 9 illustrates the steady state of the lesioned system, when probed with either “you know”, shown in panel A, or “your know”, shown in panel B. Within each panel, the top ten most active words in the first slot (left) and the second slot (right) are shown as separate nodes, and the connections linking them are shown as lines, whose width corresponds to the strength of association. Terms with higher activations are shown in darker shades of red and terms with lower activations are shown in darker shades of blue. Likewise, excitatory connections are shown in red whereas inhibitory connections are shown in blue. The line connecting “you” to “know” is dotted to indicate that it is a lesioned association. Through the global structure, activations are more inhibited for the incongruent pair compared to the congruent pair. Despite the absence of the association between the words in the grammatical pair, the greater familiarity for grammatical bigrams emerges through the mutual constraints encoded in the system.

¹⁰ Here we are using “term” instead of “word” to be more precise, because some of the items are morphemes (e.g., “nt” as in didn’t or “re” as in “you’re”).

Figure 9 shows how the scaled-up Dynamic-Eigen-Net distinguishes between grammatical (panel A: “you know”) and ungrammatical (panel B: “your know”) bigrams, even when the association forming the grammatical form (black dotted edge) has been deleted from memory. Within each panel, the top ten most active terms in each bank (first slot on the left and the second slot on the right) are indicated by separate nodes, coloured based on their relative activation. Darker shades of red correspond to more active nodes whereas darker shades of blue correspond to the less active nodes. The activations correspond to the steady-state of the system after each probe. Notably, when the probe is syntactically congruent, excitatory connections (red) outnumber inhibitory connections (blue) between the top most active terms, whereas the pattern is reversed in the case of syntactically incongruent probes.

6. Discussion

We started by asking how a system may use a limited set of exemplars to infer the structure of an infinite set that is defined by constraints on how symbolic representations within each exemplar combine. Acknowledging the immensity of the problem, we explored generalization in a very limited combinatorial domain -- the set of all well-formed bigrams. We contrasted our generalization-at-retrieval

solution to the problem with systems that attempt to generalize the structure during encoding, through error-driven learning, and proposed a modified variant of associative nets. We proposed the use of a Linear-Associative-Net equipped with dynamic connections that change the association strength between elements in the network based on the input that is processed. We provided an eigenspectrum analysis of the resulting system and described how changing weights based on input adds dynamics to the eigenspectrum of the system and expands its space of equilibria in a way that respects mutual constraints encoded in memory from prior experience.

After a series of toy demonstrations showing how the resulting system has greater generative potential than previous versions of associative nets, we turned to some empirical work using bigrams in order to explore the model's scalability and capacity to match human performance. After demonstrating the model's ability to account for basic aspects of human syntactic processing, we supplemented previous work with a bigram acceptability task. We then showed that discriminability derived from a variant of the Dynamic-Eigen-Net can predict the speed with which human participants make bigram acceptability judgements, and that such discriminability is superior to simple frequency measures. Finally, to show that the system can generalize at scale, we modeled the bigram acceptability task but deleted the associations corresponding to grammatical bigrams before probing the system. Our results showed that despite some discriminability loss, the system is able to exploit global associative structure to correctly distinguish grammatical bigrams from ungrammatical bigrams, in the absence of knowledge about each specific bigram.

In our toy demonstrations we showed how the Linear-Associative-Net suffers from severe perseveration -- it always responds with the largest eigenvector encoded in the weight matrix. We then showed how the Brain-State-in-a-Box is less vulnerable to perseveration by restricting its state-space within a hypercube that halts the gravitation of the system towards the dominant eigenvector. We then showed that the Brain-State-in-a-Box is still limited because its steady states are restricted to previously encoded patterns. In a system capable of generalization, previously encoded patterns should impose mutual constraints that enable the system to either accept or reject novel patterns depending on whether they align with the encoded structure or misalign. Our toy demonstrations of the Dynamic-Eigen-Net showed that equipping the Linear-Associative-Net with transient assemblies enables it to reach steady states to novel patterns, only if they align with the encoded structure.

In order to scale up the Dynamic-Eigen-Net, we introduced some requirements for the eigenspectrum encoded in the weight matrix. The first requirement was that the magnitude of the transient assemblies must exceed that of the largest eigenvalues of the weight matrix. The second requirement followed from the first: the rate with which the eigenvalues fall from the largest-to-the-smallest cannot be too sharp. The second requirement followed because if the eigenvalues fall too sharply, given that the magnitude of the transient assembly must exceed that of the lead eigenvalue, a sharp drop in the eigenvalues limits the state-space, such that, it is mainly influenced by the transient assembly (i.e., input to the system) and the first few eigenvectors of the system. Our way of meeting the two requirements is only provisional and may easily be accomplished in other ways. The contribution of the current manuscript is the Dynamic-Eigen-Net feedback algorithm.

We used PMI to normalize the raw co-occurrence counts. Again, we do not propose that using PMI is necessary, as other normalization measures may be used, however, it is possible that inhibitory connections are required. The role of negative evidence has been recognized in the psycholinguistic literature as necessary for the kind of generalization required for language acquisition (e.g., Pinker, 1989)¹¹. Recently, Johns et al. (2019) showed how negative information greatly improves semantic benchmarks for various representations of meaning. Our use of PMI for normalizing the weight matrix is likely benefiting from the resulting negative information. When the co-occurrence between a pair of words falls below the base-rate occurrence of each word independently, the resulting PMI yields a negative association between the two words. Our work complements work by Johns et al. (2019) by further showing how the influence of negative information can emerge through iterative feedback to drive generalization in a dynamical system.

Having scaled up the system, we started our simulations with variations of the lexical decision task. Our simulation of Goodman et al. (1981) allowed us to show how the Dynamic-Eigen-Net can track response-time costs in lexical decisions for words that are preceded by syntactically incongruent context words based on differences in familiarity of bigrams constructed using their specific stimuli. Our simulation of Colé et al. (1994) demonstrated that the system’s performance is not language-specific since we trained the model on a French corpus. Colé et al. (1994) suggested the possibility that closed-class words are distinctly represented relative to open-class words. Table 8 shows the top ten terms loading on each of the top five eigenvectors of the weight matrix. Terms that co-occur with the

¹¹ We caution readers that the “negative evidence” has been defined in various ways. We define it here as negative associations that capture anti-correlation.

highest frequency load on the top eigenvector. The second eigenvector picks up determiners, prepositions and other similar terms, in the first slot, and adjectives and nouns in the second slot. The third eigenvector has a strong affinity for pronouns, in the first slot, and verbs in the second slot. Overall, the top eigenvectors of the system have a strong propensity for capturing the variance in the use of closed-class words. Closed-class words are distinct from open-class words based on their loading on the eigenvectors of the connectivity matrix, however, such a distinction is not a function of qualitative differences in representation or encoding. The distinction between closed- and open-class words emerges from differences in their distribution in the language environment.

Table 8

Eigenvalue	First Slot	Second Slot
16764	. # to , in the " it a of	# the a of `` to it you he "
1911	the and , of # a to or `` in	small big good great natural new large young white special
903	they he i she who it we you not was	had said have made called did could found put looked
814	was is be were are been had at get but	the a his it her their them an him this
629	to can would will could may must you they should	the a it them his her their him see be

Note. The eigenvalues have been rounded to the nearest integer.

We reiterate that the current model is not meant as a complete description. The main goal of our simulations is to show that generalization is possible when the connectivity of the network is altered by the input. As pointed out by one of the reviewers, concatenation models of serial-order do not predict interference if two bigrams are encoded, with the same word (e.g., “apple”) appearing in the first slot in one bigram (“apple tree”) and in the second slot in another bigram (“juicy apple”). In one of the earliest demonstrations of same-item interference, Primoff (1938) had subjects study pairs of items in a list and later cued them with an item either in the first position or the second. The cues indicated whether the to-be-recalled paired associate was on the left or right side of the cue word. Primoff (1938) found superior cued recall for items that only appeared once in a list relative to items that appeared twice -- once as the first member of a pair and another time as the second member of a different pair. Recently,

Rehani and Caplan (2011) replicated the findings and suggested that they require modifications to models that encode serial-position through concatenation, as we have done here.

We may be able to accommodate the result based on the interaction between serial-order and syntagmatic associations. Syntagmatic associations correspond to order-independent links between words that occur in the same context. The simulations presented here only incorporated the former kind of association, however, we believe that a more complete model will require the addition of the latter. Figure 11 shows one way the two kinds of associations may interact. We can encode serial-order associations in the first two slots as we have done here, but add a third slot to include syntagmatic associations. For instance, when a subject studies the pair A-B, the unit for A would be turned on (e.g., set to one) in the first slot, the unit for B would be turned on in the second slot, and both units A and B would be turned on in the third slot. The outer-product of the vector with itself would encode forward associations in the top-middle submatrix and backward associations in the middle-left submatrix. The lower-right submatrix would encode syntagmatic associations (i.e., many-to-many links) and the two submatrices above and the two submatrices to the left would encode item associations (i.e., one-to-many links). Murdock (1974) provides an in-depth examination of order and item information, and Dennis (2005) shows how syntagmatic and order-based relations can be combined to capture several problems in semantic composition. If we assume another pair, B-C that was encoded, probing the system with B (i.e., turning on the unit for B in the first and third slots) would activate both associates A-B and B-C via spreading activation through the item and syntagmatic associations, leading to interference (*cf* Kato & Caplan, 2017).

The diagram illustrates the interaction of serial-order information with item and syntagmatic information. It consists of a 3x3 grid of associations and an elongated rectangle above it representing the input sequence A-B.

The elongated rectangle is divided into three segments: A, B, and A B. The 3x3 grid is divided into three rows and three columns. The rows are labeled A, B, and B A on the left. The columns are labeled A, B, and A B on top.

	A	B	A B
A		FW	ITEM
B	BW		ITEM
B A	ITEM	ITEM	SYN

Figure 11 shows one way serial-order information can interact with item and syntagmatic information. The elongated rectangle above the square partitioned table corresponds to the pair A-B. The item, A, is active in the first slot and the item, B, is active in the second slot. Both A and B are active in the third slot. The outer-product of such input with itself yields a block-matrix that can be partitioned into forward (FW), backward (BW), item (ITEM), and syntagmatic (SYN) associations.

An open question remains as to how serial-order and syntagmatic associations may interface. The experiments conducted by Goodman et al. (1981) and Münte et al. (1993) manipulated both syntactic and semantic congruity in their tasks, however, we have only explored syntactic congruity in the current work in order to demonstrate structural, or combinatorial, generalization. One way semantic congruity is explored is in lexical decision tasks where the second word in two adjacent word-present trials is synonymous (e.g., “feline cat”) with the first word or when it is associatively related (“cat fluffy”). Some of the work exploring the interaction between syntactic and semantic congruity (e.g., Goodman et al., 1981; Münte et al., 1993; Seidenberg et al., 1994) has led to the suggestion that syntactic and semantic primes tap into different processing loci, with the overall theme being that semantic effects are stronger, automatic, and occur earlier in processing while syntactic effects are weaker, controlled, and occur later in processing. A major challenge has been to orthogonally manipulate both semantic and syntactic variables, as they are typically highly confounded, i.e., syntactically anomalous utterances are also often semantically anomalous. Our current view is that associations of different kinds underlie semantic and syntactic effects, and that the interactions between the two occur within a highly interactive system. The addition of syntagmatic associations to serial-order associations would be a natural path forward for capturing some of the semantic-syntactic interactions documented in the literature.

The highly recurrent architecture of associative nets was originally motivated by the observation that cortical pyramidal cells are subsumed in a vast network of feedback connections (e.g., Douglas et al., 1995). Similarly, Hebbian encoding has long been considered the computational analogue to Long-Term Potentiation (Collingridge & Bliss, 1995). In the same way, transient weights in the Dynamic-Eigen-Net may be grounded in synaptic plasticity that spans time-scales ranging from hundreds of milliseconds to a few minutes, referred to generally as Short-Term-Plasticity¹² (Masse et al., 2020). Zucker & Regehr, 2002 review work around Short-Term-Plasticity and suggest evidence favouring the “residual Ca^{2+} hypothesis”. The hypothesis is that Ca^{2+} ions linger at the presynaptic site following an action-potential and keep the synapse more excitable relative to rest. Much of the theoretical work exploring Short-Term-Plasticity has used spiking neuron models, where action-potentials are modeled directly by specifying single neuron dynamics (e.g., using the integrate-and-fire model). What the spiking neuron models have in common with the Dynamic-Eigen-Net is that both conceptualize neural computation using attractor dynamics in an associative net (e.g., Lundqvist et al., 2006). An important difference, however, is that in our modeling the values in the state-vector correspond to rates of firing and not the cell membrane potential, therefore allowing us to abstract over cellular dynamics and work with a linear system (see Anderson & Silverstein’s, 1978, response to Grossberg, 1978 for an example of why confusing the two can be theoretically problematic). In Marr’s terms, the spiking neuron models are closer to the implementation level whereas our model is closer to the algorithmic level -- much can be gained by bridging the gap.

Although a systematic exploration of Short-Term-Plasticity in associative nets has not been presented in the recent cognitive modelling literature, its use in models has precedence (e.g., Gardner-Medwin, 1989; Burgess & Hitch, 1999; Plaut & Shallice, 1993; Feldman, 1982). The various implementations make different architectural and learning assumptions, but generally show that including Short-Term-Plasticity has several advantages over networks with static connectivity. For instance, in the Hebbian systems, Gardner-Medwin (1989) assumes short-term weights are multiplicatively combined with long-term weights whereas Burgess and Hitch (1999) assume the additive combination of the two. Gardner-Medwin (1989) provides information-theoretic analyses

¹² We use the term Short-Term-Plasticity only to mean changes to synaptic conductances that span seconds to minutes as opposed to long-term changes that span over days to months. To our knowledge, except for the spiking neuron modeling, implementations of Short-Term-Plasticity in more abstract models do not explicitly specify temporal dynamics, as may be specified through a time-constant in a differential equation.

showing capacity advantages using their implementation of Short-Term-Plasticity and Burgess and Hitch (1999) rely on Short-Term-Plasticity to capture a set of memory benchmarks. Plaut and Shallice (1993) additively combine Short-Term-Plasticity to long-term weights, trained using error-driven learning, to model perseveration in their model of deep dyslexia. Feldman (1982) suggests Short-Term-Plasticity as one possible mechanism for feature-binding in vision.

Short-term weights have also been explored in the broader machine-learning literature, however, the implementations rely on error-driven learning. For instance, Ba, Hinton, Mnih, Leibo, & Ionesco (2016) suggest encoding a set of recent hidden states into a fast-decaying set of short-term weights as a way to model greater facilitation for states corresponding to the system's recent hidden state history. They propose the mechanism for recurrent neural nets which use error-driven learning for training the long-term connectivities. Reliance on error-driven learning limits the system's ability to meet human-level cognitive benchmarks such as one-shot learning. Perhaps the kind of generalization sought in the broad machine-learning literature is better driven through recurrence within a Dynamic-Eigen-Net as opposed to slow error-driven learning, at least in some domains.

In conclusion, our demonstrations show that including transient assemblies in Linear-Associative-Nets increases their generative complexity. Given that generalization requires exploiting structural regularities with little surface-level information, and a Dynamic-Eigen-Net's ability to assimilate novel patterns based on structure alone, we propose our model as one candidate for a system capable of structural generalization at retrieval.

References

- Amari, S. I. (1977). Neural theory of association and concept-formation. *Biological Cybernetics*, 26(3), 175-185.
- Anderson, J. A., Silverstein, J. W., Ritz, S. A., & Jones, R. S. (1977). Distinctive features, categorical perception, and probability learning: Some applications of a neural model. *Psychological Review*, 84(5), 413-451.
- Anderson, J. A. (1995). An introduction to neural networks. *MIT press*.
- Ba, J., Hinton, G., Mnih, V., Leibo, J. Z. & Ionesco, C. (2016). Using fast weights to attend to the recent past. *Advances in Neural Information Processing Systems*, 29, 4331-4339
- Barnes, J. M., & Underwood, B. J. (1959). Fate" of first-list associations in transfer theory. *Experimental Journal of Psychology*, 58, 97-105.
- Burgess, N. & Hitch, G. J. (1999). Memory for serial order: a network model of the phonological loop and its timing. *Psychological Review*, 106(3), 551-581.
- Collingridge, G. L., Bliss, T. V. P. (1995). Memories of NMDA receptors and LTP. *Trends in Neurosciences*, 18, 54-56.
- Cole, P., & Segui, J. (1994). Grammatical incongruency and vocabulary types. *Memory & Cognition*, 22(4), 387-394.
- Chubala, C. M., & Jamieson, R. K. (2013). Recoding and representation in artificial grammar learning. *Behavior Research Methods*, 45(2), 470-479.
- Church, K., & Hanks, P. (1990). Word association norms, mutual information, and lexicography. *Computational linguistics*, 16(1), 22-29.
- Dennis, S. (2005). A memory-based theory of verbal cognition. *Cognitive Science*, 29(2), 145-193.
- Douglas, R. J., Koch, C., Mahowald, M., Martin, K. A., & Suarez, H. H. (1995). Recurrent excitation in neocortical circuits. *Science*, 269(5226), 981-985.
- Farrell, S., & Lewandowsky, S. (2002). An endogenous distributed model of ordering in serial recall. *Psychonomic bulletin & review*, 9(1), 59-79.
- Feldman, J. A. (1982). Dynamic connections in neural networks. *Biological Cybernetics*, 46, 27-39.
- Gardner-Medwin, A. R. (1989). Doubly modifiable synapses: a model of short and long term auto-associative memory. *Proc. R. Soc. Lond. B* 238, 137-154.

- Garrett, M. F. (1978). Word and sentence perception. In *Perception* (pp. 611-625). Springer, Berlin, Heidelberg.
- Grefenstette, G. (1994). Corpus-derived First, Second, and Third-order Word Affinities (pp. 279-290). *Rank Xerox Research Centre*.
- Goodman, G. O., McClelland, J. L., & Gibbs, R. W. (1981). The role of syntactic context in word recognition. *Memory & Cognition*, 9(6), 580-586
- Hebb, D. O. *The organization of behavior*. New York: Wiley, 1949.
- Hinton, G. E. (1990). Mapping part-whole hierarchies into connectionist networks. *Artificial Intelligence*, 46(1-2), 47-75.
- Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological review*, 93(4), 411.
- Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological review*, 95(4), 528.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. USA*, 79, 2554-2558.
- Honnibal, M., Montani, I., Van Landeghem, S., & Boyd, A. (2016). spacy: Industrial-strength natural language processing in python. *spaCy*. <https://spacy.io/>.
- Jamieson, R. K., & Mewhort, D. J. K. (2009). Applying an exemplar model to the artificial-grammar task: Inferring grammaticality from similarity. *Quarterly Journal of Experimental Psychology*, 62(3), 550-575.
- Jamieson, R. K., & Mewhort, D. J. K. (2010). Applying an exemplar model to the artificial-grammar task: String completion and performance on individual items. *Quarterly Journal of Experimental Psychology*, 63(5), 1014-1039.
- Jamieson, R. K., & Mewhort, D. J. K. (2011). Grammaticality is inferred from global similarity: A reply to Kinder (2010). *Quarterly Journal of Experimental Psychology*, 64(2), 209-216.
- Johns, B. T., & Jones, M. N. (2015). Generating structure from experience: A retrieval-based model of language processing. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, 69(3), 233
- Johns, B. T., Mewhort, D. J., & Jones, M. N. (2019). The role of negative information in distributional semantic learning. *Cognitive Science*, 43(5), e12730.

- Johns, B. T., Jamieson, R. K., Crump, M. J., Jones, M. N., & Mewhort, D. J. K. (2020). Production without rules: Using an instance memory model to exploit structure in natural language. *Journal of Memory and Language*, 115, 104165.
- Jones, M. N., & Mewhort, D. J. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological review*, 114(1), 1.
- Jones, M. N. (2019). When does abstraction occur in semantic memory: insights from distributional models. *Language, Cognition and Neuroscience*, 34(10), 1338-1346.
- De Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior research methods*, 47(1), 1-12.
- Kato, K., & Caplan, J. B. (2017). Order of items within associations. *Journal of Memory and Language*, 97, 81-102.
- Kintsch, W., & Walter Kintsch, (1998). Comprehension: A paradigm for cognition. *Cambridge university press*.
- Kwantes, P. J. (2005). Using context to build semantics. *Psychonomic Bulletin & Review*, 12(4), 703-710.
- Mannering, W., & Jones, M. N. (2020). Catastrophic Interference in Predictive Neural Network Models of Distributional Semantics. *Computational Brain & Behavior*.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological review*, 88(5), 375.
- McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation (Vol. 24, pp. 109-165)*. Academic Press.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems (pp. 3111-3119)*.
- Murdock, B. B. (1974). Human memory: Theory and data. Lawrence Erlbaum.
- Münste, T. F., Heinze, H. J., & Mangun, G. R. (1993). Dissociation of brain activity related to syntactic and semantic aspects of language. *Journal of Cognitive Neuroscience*, 5(3), 335-344.
- Masse, N. Y., Rosen, M. C., & Freedman, D. J. (2020). Reevaluating the role of persistent neural activity in short-term memory. *Trends in cognitive sciences*, 24(3), 242-258.

- Plaut, D. C. & Shallice, T. Perseverative and semantic influences on visual object naming errors in optic aphasia: a connectionist account. *Journal of Cognitive Neuroscience*, 5(1), 89-117.
- Rogers, T. T., & McClelland, J. L. (2014). Parallel distributed processing at 25: Further explorations in the microstructure of cognition. *Cognitive science*, 38(6), 1024-1077.
- Pinker, S. (1989). Learnability and cognition: The acquisition of argument structure.
- Primoff, E. (1938). Backward and forward association as an organizing act in serial and in paired associate learning. *The Journal of Psychology*, 5(2), 375-395.
- Rehani, M., & Caplan, J. B. (2011). Interference and the representation of order within associations. *Quarterly Journal of Experimental Psychology*, 64(7), 1409-1429.
- Rumelhart, D. E., & McClelland, J. L. (1986). On learning the past tenses of English verbs.
- Sahlgren, M., Holst, A., & Kanerva, P. (2008). Permutations as a means to encode order in word space. *In The 30th Annual Meeting of the Cognitive Science Society (CogSci'08), 23-26 July 2008, Washington DC, USA.*
- Seidenberg, M. S., Waters, G. S., Sanders, M., & Langer, P. (1984). Pre- and postlexical loci of contextual effects on word recognition. *Memory & cognition*, 12(4), 315-328.
- Smolensky, P. (1986). Information Processing in Dynamical Systems: Foundations of Harmony Theory. *Parallel distributed processing: explorations in the microstructure of cognition*, 1.
- Stickgold, R., & Walker, M. P. (2013). Sleep-dependent memory triage: evolving generalization through selective processing. *Nature neuroscience*, 16(2), 139.
- Toutanova, K., Klein, D., Manning, C., & Singer, Y. (2003). Feature-Rich Part-of-Speech Tagging with a Cyclic Dependency Network. *Proc. HLT-NAACL 2003*, pp, 252-259.
- Westbury, C., Hollis, G. (2019). Conceptualizing syntactic categories as semantic categories: Unifying part-of-speech identification and semantics using co-occurrence vector averaging. *Behavior Research Methods*, 51, 1371-1398.
- Zucker, R. S. & Regehr, W. G. (2002). Short-term synaptic plasticity. *Annual Review of Physiology*, 64, 355-405.