

Resonator Networks, 1: An Efficient Solution for Factoring High-Dimensional, Distributed Representations of Data Structures

E. Paxon Frady

epaxon@berkeley.edu

Redwood Center for Theoretical Neuroscience, University of California, Berkeley, Berkeley, CA 94720, U.S.A., and Intel Laboratories, Neuromorphic Computing Lab, San Francisco, CA, 94111, U.S.A.

Spencer J. Kent

spencer.kent@berkeley.edu

Redwood Center for Theoretical Neuroscience and Electrical Engineering and Computer Sciences, University of California, Berkeley, Berkeley, CA 94720, U.S.A.

Bruno A. Olshausen

baolshausen@berkeley.edu

Redwood Center for Theoretical Neuroscience, Helen Wills Neuroscience Institute, and School of Optometry, University of California, Berkeley, Berkeley, CA 94720, U.S.A.

Friedrich T. Sommer

fsommer@berkeley.edu

Redwood Center for Theoretical Neuroscience and Helen Wills Neuroscience Institute, University of California, Berkeley, Berkeley, CA 94720, U.S.A., and Intel Laboratories, Neuromorphic Computing Lab, San Francisco, CA 94111, U.S.A.

The ability to encode and manipulate data structures with distributed neural representations could qualitatively enhance the capabilities of traditional neural networks by supporting rule-based symbolic reasoning, a central property of cognition. Here we show how this may be accomplished within the framework of Vector Symbolic Architectures (VSAs) (Plate, 1991; Gayler, 1998; Kanerva, 1996), whereby data structures are encoded by combining high-dimensional vectors with operations that together form an algebra on the space of distributed representations. In particular, we propose an efficient solution to a hard combinatorial search problem that arises when decoding elements of a VSA data structure: the factorization of products of multiple codevectors. Our proposed algorithm, called a resonator network, is a new type of recurrent neural network that interleaves VSA multiplication operations and pattern completion. We show in two examples—parsing of a tree-like data

structure and parsing of a visual scene—how the factorization problem arises and how the resonator network can solve it. More broadly, resonator networks open the possibility of applying VSAs to myriad artificial intelligence problems in real-world domains. The companion article in this issue (Kent, Frady, Sommer, & Olshausen, 2020) presents a rigorous analysis and evaluation of the performance of resonator networks, showing it outperforms alternative approaches.

1 Introduction

Cognition requires making use of learned knowledge in contexts never before encountered, a facility that requires information to be represented in terms of components that may be flexibly recombined. A long-standing goal for neuroscience and psychology has been to understand how such capacities are expressed by neural networks in the brain. Early artificial intelligence researchers developed frameworks of symbol manipulation to emulate cognition, but they were implemented with local data representations (where the meaning of a bit is tied to its location) that are brittle and nonadaptive (Kanerva, 1997). Connectionism, a movement started in psychology (McClelland, Rumelhart, & PDP Research Group, 1986), based itself on the premise that internal representations of knowledge must be highly distributed and be able to adapt to the statistics of the data so as to learn by example. Along the way, however, connectionism also gave up many of the rich capabilities offered by symbolic computation (Jackendoff, 2002). In recent years, it has become clear that a unification of the ideas behind each approach—distributed representation, adaptivity, and symbolic manipulation—will be required for reproducing the brain’s ability to learn from few examples, to deal with novel situations, or to change behaviors when driven by internal information processing rather than purely by external events (Plate, 2003; Gayler, 2003; Kanerva, 2009; Lake, Ullman, Tenenbaum, & Gershman, 2017).

Digital computers owe their power and ubiquity to the abstraction of data structures, which support decomposing information into parts, referencing each part individually, and composing these parts with other data structures. Examples include trees, records with fields, and linked lists. Connectionist theories have long been criticized because it is hard to imagine how compound, hierarchical data structures could be represented and manipulated by neural networks (Hinton, 1990). Cognitive scientists have argued that at the very least, cognitive data structures should support three patterns of combination, which are familiar to any computer programmer (Fodor & Pylyshyn, 1988):

1. Key-value pairs: A key or variable is a placeholder for information to which a value can be assigned in a particular instance. This

association, variable binding, generates what is called the system-aticity of cognition (Fodor, 1975; Plate, 2003).

2. Sequential structures: A *sequence* is an ordered pattern of organization and computation required by many reasoning tasks.
3. Hierarchy: The notion that some aspects of knowledge can be decomposed recursively into a set of successively more fundamental parts.

Variable binding, sequence, and hierarchy are critical structures of cognition, and a comprehensive theory of intelligence must take these into account.

A family of models called Vector Symbolic Architectures (VSAs) encodes these structures into distributed representations, providing a framework that can reconcile the symbolic and connectionist perspectives (Plate, 2003; Gayler, 2003; Kanerva, 2009). Building on the concept of reduced representations (Hinton, 1990), VSAs allow one to express data structures holographically in a vector space of high but fixed dimensionality. The atoms of representation are random high-dimensional vectors, and data structures built from these atoms are vectors with the same dimension. Three operations are used to form and manipulate data structures—addition, multiplication, and permutation—which together form an algebra over the space of high-dimensional vectors. These operations enable building representations of sets, ordered lists (sequences), n-tuples, trees, key-value bindings, and records containing role-filler relationships which can be composed into hierarchies, as described in Plate (1995); Kanerva (1996, 1997); Joshi, Halseth, and Kanerva (2016); Frady, Kleyko, and Sommer (2018), and below.

In order to read out or access the components of a VSA-encoded data structure, the high-dimensional vector representing it must be decomposed into the primitives or atomic vectors from which it is built. This is the problem of decoding. For example, if the primitives are combined by addition only, the distributed representation can be decoded by a nearest-neighbor look-up or an autoassociative memory. However, hierarchical or compound data structures, such as a multilevel tree or an object with multiple attributes bound together, are built from combinations of addition, multiplication, and permutation operations on the primitives. In this case, decoding via a simple nearest-neighbor look-up would require storing every possible combination of the primitives (e.g., all possible paths in a tree or all the possible attribute combinations) essentially amounting to a combinatoric search problem. Past applications of VSAs have largely sidestepped this problem by limiting the depth of the data structures or using a brute force approach to consider all possible combinations when necessary (Plate, 2000a; Cox, Kachergis, Recchia, & Jones, 2011). As a result, the application of VSAs to real-world problems has been rather limited, since up to now, there has not been a solution for efficiently accessing elements of such compound data structures containing a product of multiple components.

The solution to this dilemma is to factorize the high-dimensional vector representing a compound data structure into the primitives from which it is composed. That is, given a high-dimensional vector formed from an element-wise product of two or more vectors, we must find its factors. This way, a nearest-neighbor look-up need only search over the alternatives for each factor individually rather than all possible combinations. Obviously, though, factorization poses a difficult computational problem in its own right.

Here, we propose an efficient algorithm for factorizing high-dimensional vectors that may be interpreted as a type of recurrent neural network, which we call a *resonator network*. The resonator network relies on the VSA principle of superposition to search through the combinatoric solution space without directly enumerating all possible factorizations. Given a high-dimensional vector as input, the network iteratively searches through many potential factorizations in parallel until a set of factors is found that agrees with the input. Solutions emerge as stable fixed points in the network dynamics.

In this article, part 1 of a two-part series in this issue, we first briefly introduce the VSA framework and the problem of factoring high-dimensional VSA representations. We then show using two examples—searching a binary tree and querying the contents of a visual scene—how VSAs may be used to build distributed representations of compound data structures and how resonator networks are used to decompose these data structures and solve the problem. The companion article in this issue (Kent, Frady, Sommer, & Olshausen, 2020) provides rigorous mathematical and simulation analysis of resonator networks and compares its performance with alternative approaches for solving high-dimensional vector factorization problems.

2 VSA Preliminaries

All entities in a VSA are represented as high-dimensional vectors in the same space, with vector dimension N typically in the range of 1000 to 10,000. In this article, we focus on the VSA framework called Multiply-Add-Permute (Gayler, 1998, 2003). The atomic primitives are bipolar vectors whose components are ± 1 , chosen randomly. These vectors are used as symbols to represent concepts. The set of atomic vectors representing specific items is stored in a codebook, which is a matrix of dimension $N \times D$, where D is the number of atoms.

The use of high-dimensional vectors is an important aspect of the VSA framework, as it relies on the concentration of measure phenomenon (Ledoux, 2001) that independently chosen random vectors are very close to orthogonal, a property we refer to as *quasi-orthogonality*. This property allows vectors to act symbolically, as the similarity (inner product) between two different atomic vectors is small compared to their self-similarity

(L2 norm). Furthermore, a much larger set of quasi-orthogonal vectors exists than orthogonal vectors, which may be exploited for combinatoric search.

Data structures are composed and computations are carried out via an algebra consisting of three vector operations: addition, multiplication, and permutation. The elements of a data structure are then read out (decoded) using the conventional vector dot product as a similarity measure to compare to items stored in the codebook. The VSA operations of addition, multiplication, and permutation act to manipulate the vector symbols in ways that preserve or destroy their similarity.

Formally, the VSA operations are defined as follows:

Dot product (\cdot) is the conventional vector inner product, $\mathbf{x} \cdot \mathbf{s} = \sum_i x_i s_i$, which is used to measure the similarity between vectors. This is used to decode the result of a VSA computation by comparing the vector to the set of vectors in the codebook:

$$\mathbf{a} = \mathbf{X}^\top \mathbf{s}.$$

Here, \mathbf{X} is the codebook of atomic vectors, and \mathbf{s} is a high-dimensional vector resulting from a VSA computation. The result of a VSA computation can be a single symbol indicated by the largest component of \mathbf{a} . Alternatively, the coefficients \mathbf{a} can be considered as a weighted sum, where each entry indicates a confidence level, probability, or intensity value.

Addition ($+$) is used to superpose items together, like forming a set. It is defined by regular vector addition, the element-wise sum:

$$\mathbf{s} = \mathbf{x} + \mathbf{y},$$

or $s_i = x_i + y_i$. Depending on the circumstances, the sum may be kept as is or subsequently thresholded so that each s_i is ± 1 . In either case, the addition operation results in a vector that is similar to each of its superposed components; one can determine the members of the sum by similarity to the atomic vectors. Superposition is possible because of the quasi-orthogonal property. However, superposition produces a small amount of cross-talk noise, which increases with the number of items in the sum and is diminished with large vector dimensionality (see Frady et al., 2018, for a detailed characterization of superposition).

Multiplication (\odot) is used to bind items together to form a conjunction, such as in assigning a value to a variable. It is defined by the Hadamard product between vectors, that is, the element-wise multiplication of vector components:

$$\mathbf{s} = \mathbf{x} \odot \mathbf{y},$$

or $s_i = x_i y_i$. This multiplication operation is invertible ($\mathbf{y} = \mathbf{s} \odot \mathbf{x}$), and it distributes over addition, $\mathbf{x} \odot \mathbf{y} + \mathbf{x} \odot \mathbf{z} = \mathbf{x} \odot (\mathbf{y} + \mathbf{z})$. Note that in the MAP VSA, the bipolar primitive vectors are their own self-inverses. In contrast to addition, multiplication generates a vector that is dissimilar to each of its inputs (Kanerva, 2009).

Permutation ($\rho(\cdot)$) is used to “protect” or “order” items. It operates on a single input vector. In principle, it can be any random permutation, but is typically a simple cyclic shift:

$$\mathbf{s} = \rho(\mathbf{x}),$$

or $s_i = x_{(i-1)\%N}$. Permutation distributes over both addition, $\rho(\mathbf{x}) + \rho(\mathbf{y}) = \rho(\mathbf{x} + \mathbf{y})$, and multiplication, $\rho(\mathbf{x}) \odot \rho(\mathbf{y}) = \rho(\mathbf{x} \odot \mathbf{y})$, and its function is complementary to addition and multiplication. Permutations are used to protect the components of a data structure built with these other operations, based on the fact that permutation and binding are noncommutative, $\mathbf{x} \odot \rho(\mathbf{y}) \neq \mathbf{y} \odot \rho(\mathbf{x})$. In essence, permutation rotates vectors into dimensions of the space that are almost orthogonal to the dimensions used by the original vectors. Information is thus protected when combined with other items, because vector components will not appear similar to or interfere with those other items. Permutations can also be used to index sequences (Frady et al., 2018), or levels in a hierarchy, by successive application of the permutation operation. For example to represent the sequence $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2$ in a vector $\mathbf{s} = \mathbf{x}_0 + \rho(\mathbf{x}_1) + \rho^2(\mathbf{x}_2)$, with $\rho^2(\mathbf{x}) = \rho(\rho(\mathbf{x}))$.

VSA's combine these operations to form data structures and to compute with them. The combination of atomic vectors into composite data structures is rather straightforward. But as we shall see, querying composite data structures often results in the problem of decoding terms composed of two (or perhaps many more) atomic vectors that are multiplied together. In order to decode such composite vectors, one must search through many combinations of atoms. In general, this is a hard combinatorial search problem, which typically requires directly testing every combination of factors. The resonator network can efficiently solve these problems without needing to directly test every combination of factors.

3 Factorization via Search in Superposition

In general, the factorization problems that arise in VSA's may involve two or more factors, but let us assume we are given a composite vector \mathbf{s} , formed as a product of three vectors,

$$\mathbf{s} = \mathbf{x}_i^* \odot \mathbf{y}_j^* \odot \mathbf{z}_k^*, \quad (3.1)$$

where the vectors \mathbf{x}_{i^*} , \mathbf{y}_{j^*} , and \mathbf{z}_{k^*} are drawn from codebooks $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_D\}$, $\mathbf{Y} = \{\mathbf{y}_1, \dots, \mathbf{y}_D\}$, and $\mathbf{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_D\}$. Given \mathbf{s} and the codebooks \mathbf{X} , \mathbf{Y} , and \mathbf{Z} , the task is to find \mathbf{x}_{i^*} , \mathbf{y}_{j^*} , and \mathbf{z}_{k^*} .

The resonator network is an iterative approach to solve this problem without exhaustively searching through each possible combination of the factors. A key motivating idea behind resonator networks is the VSA principle of superposition. In VSAs, multiple symbols can be expressed simultaneously in a single high-dimensional vector via vector addition. Randomized atomic vectors are highly likely to be close to orthogonal in high-dimensional space, meaning that they can be superposed without much interference. However, there is some cross-talk noise between the superposed symbols, and “clean-up memory” (such as a Hopfield network) is thus utilized to reduce the cross-talk noise.

A resonator network combines the strategy of superposition and cleanup memory to efficiently search over the combinatorially large space of possible factorizations. The vectors $\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$, and $\hat{\mathbf{z}}$ represent the current estimate for each factor. These vectors can be initialized to the superposition of all possible factors—for example, $\hat{\mathbf{x}}(0) = \sum_i^D \mathbf{x}_i$, $\hat{\mathbf{y}}(0) = \sum_j^D \mathbf{y}_j$. A particular factor can then be inferred from \mathbf{s} based on the estimates for the other two—for example, $\hat{\mathbf{z}}(1) = \mathbf{s} \odot \hat{\mathbf{x}}(0) \odot \hat{\mathbf{y}}(0)$. Since binding distributes over addition, the product $\hat{\mathbf{x}}(0) \odot \hat{\mathbf{y}}(0)$ expresses every combination of factors in superposition because $\hat{\mathbf{x}}(0) \odot \hat{\mathbf{y}}(0) = \sum_i^D \sum_j^D \mathbf{x}_i \odot \mathbf{y}_j$. For instance, if $D = 100$, then this initial guess represents $D^2 = 10,000$ combinations in superposition. Thus, many potential combinations of the pair of factors may be considered at once when inferring the third factor.

The inference process, however, is noisy if many guesses are tested simultaneously. This noise results from cross talk of many quasi-orthogonal vectors and can be reduced through a clean-up memory. This is built from the codebooks, which contain all the vectors that are possible factors of the input \mathbf{s} . Each clean-up memory projects the initial noisy estimate onto the span of the codebook. This computes a measure of confidence for whether each element in the codebook is a factor.

The result of the inference and clean-up leads to a new estimate for each factor. The new estimate is formed by a sum of dictionary items weighted by the confidence levels. This produces a better guess for each one of the factors. The inference can then be repeated with better guesses, which reduces cross-talk noise even further. By iteratively applying this procedure, the inference and clean-up stages cooperate to successively reduce cross-talk noise until the solution is found.

The procedure described above, for all three factors, is specified by the following set of equations (see Figure 1):

$$\begin{aligned}\hat{\mathbf{x}}(t+1) &= g(\mathbf{X}\mathbf{X}^\top (\mathbf{s} \odot \hat{\mathbf{y}}(t) \odot \hat{\mathbf{z}}(t))), \\ \hat{\mathbf{y}}(t+1) &= g(\mathbf{Y}\mathbf{Y}^\top (\mathbf{s} \odot \hat{\mathbf{x}}(t) \odot \hat{\mathbf{z}}(t))), \\ \hat{\mathbf{z}}(t+1) &= g(\mathbf{Z}\mathbf{Z}^\top (\mathbf{s} \odot \hat{\mathbf{x}}(t) \odot \hat{\mathbf{y}}(t))),\end{aligned}\tag{3.2}$$




Figure 1: A resonator network with three factors.

where the function g prevents runaway positive feedback by thresholding the elements of each vector to ± 1 .

If we examine the clean-up memory for \hat{x} , which contains a matrix multiplication with $\mathbf{X}\mathbf{X}^\top$ and thresholding function g , then we see this operation is nearly identical to a Hopfield network with outer-product Hebbian learning (Hopfield, 1982). Except here, rather than directly feeding back into itself, the result of the clean-up is sent to other parts of the network.

The set of equations in (3.2) defines a nonlinear dynamical system that has interesting empirical and theoretical properties, which we thoroughly examine through simulation experiments in Kent et al. (2020), the companion article in this issue. Empirically, the system bounces around in state space until the correct solution appears to resonate with the network dynamics, popping out as if in a moment of insight. We find that while there is no Lyapunov function governing these dynamics and no guarantee for convergence, the resonator network empirically converges to the correct solution with high probability as long as the number of product combinations to be searched is within the network’s operational capacity. We show that the operational capacity is given by a quadratic function of N . Compared to numerous alternative optimization methods that we considered, this capacity for resonator networks is higher by almost two orders of magnitude.

4 Decoding Data Structures with Resonator Networks

We now turn to two examples that illustrate how VSA operations can be combined to build distributed representations of data structures, how the factorization problem arises when parsing these representations, and how resonator networks can be designed to solve this problem.

4.1 Searching a Tree Data Structure. Consider the tree data structure depicted in Figure 2. We can form a distributed representation of this tree in a single high-dimensional vector by using all three VSA operations: superposition $+$, binding \odot , and permutation $\rho(\cdot)$. First, each leaf in the tree is assigned a random vector $\mathbf{a}, \mathbf{b}, \dots, \mathbf{g} \in \{-1, +1\}^N$. We also assign random

vectors **left** and **right** that are used to describe position in the tree. Moving from the root of the tree to a particular leaf involves a sequence of **left** and **right** turns. The order of these turns is represented by permutation $\rho(\cdot)$. The number of times permutation is applied indicates depth within the tree: **left** is a **left** turn at depth 0, $\rho(\mathbf{left})$ is a **left** turn at depth 1, $\rho^2(\mathbf{left})$ is a **left** turn at depth 2, and so on. A sequence of turns is represented by the binding of these vectors; for example, $\mathbf{left} \odot \rho(\mathbf{left}) \odot \rho^2(\mathbf{left})$ corresponds to three left turns. We can then attach to each leaf its position in the tree, again with binding: for example, $\mathbf{a} \odot \mathbf{left} \odot \rho(\mathbf{left}) \odot \rho^2(\mathbf{left})$. Finally, the representation for the whole tree is collapsed into a single vector, **tree**, via superposition:

$$\begin{aligned}
 \mathbf{tree} = & \mathbf{a} \odot \mathbf{left} \odot \rho(\mathbf{left}) \odot \rho^2(\mathbf{left}) \\
 & + \mathbf{b} \odot \mathbf{left} \odot \rho(\mathbf{right}) \odot \rho^2(\mathbf{left}) \\
 & + \mathbf{c} \odot \mathbf{right} \odot \rho(\mathbf{right}) \odot \rho^2(\mathbf{left}) \\
 & + \mathbf{d} \odot \mathbf{right} \odot \rho(\mathbf{right}) \odot \rho^2(\mathbf{right}) \odot \rho^3(\mathbf{left}) \\
 & + \mathbf{e} \odot \mathbf{right} \odot \rho(\mathbf{right}) \odot \rho^2(\mathbf{right}) \odot \rho^3(\mathbf{right}) \\
 & + \mathbf{f} \odot \mathbf{left} \odot \rho(\mathbf{right}) \odot \rho^2(\mathbf{right}) \odot \rho^3(\mathbf{left}) \odot \rho^4(\mathbf{left}) \\
 & + \mathbf{g} \odot \mathbf{left} \odot \rho(\mathbf{right}) \odot \rho^2(\mathbf{right}) \odot \rho^3(\mathbf{left}) \odot \rho^4(\mathbf{right}). \quad (4.1)
 \end{aligned}$$

The vector **tree** encodes the information so that we can flexibly query the data structure using VSA operations. For instance, we can find the identity of the leaf located at position **left**, **right**, **left** by “unbinding” the representation of this location from the vector representing the tree. Binding and unbinding are performed with the same operation since bipolar vectors are self-inverses. When we unbind the query location by Hadamard product, it will distribute through the superposition and cancel out with itself, leaving the atomic vector attached to that location “exposed”:

$$\mathbf{tree} \odot (\mathbf{left} \odot \rho(\mathbf{right}) \odot \rho^2(\mathbf{left})) = \mathbf{b} + \mathit{noise}. \quad (4.2)$$

The noise term arises since the query distributes through the sum. The other terms combine with the query but remain quasi-orthogonal to the vectors stored in the codebook, which keeps the other items in the tree “hidden.” That is, the terms contained in *noise* are dissimilar from each of the atoms stored in the codebook, and this appears as gaussian noise when decoding (Fraday et al., 2018). The vector $\mathbf{b} + \mathit{noise}$ will have high similarity with atom **b** in the codebook and will be successfully decoded by nearest-neighbor or associative memory lookup among the atoms with high probability.

With this flexible encoding of the data structure, instead of asking for the label at a specific position, we can ask for the position of a specific label

(essentially the problem of tree search). For instance, the query that exposes the position of label c is simply

$$\mathbf{tree} \odot \mathbf{c} = \mathbf{right} \odot \rho(\mathbf{right}) \odot \rho^2(\mathbf{left}) + \mathit{noise}. \quad (4.3)$$

This presents a new challenge, however, because we still need to decode the composite vector $\mathbf{right} \odot \rho(\mathbf{right}) \odot \rho^2(\mathbf{left}) + \mathit{noise}$ into the parts that describe a position in the tree. In previous applications of VSAs, one would exhaustively enumerate all traversals of the tree and compute similarity to find the path. Instead, we can use a resonator network.

To set up the network for this problem, we first establish a maximum depth to search through; the maximum depth determines the number of factors that need to be estimated. For the tree shown in Figure 2, we need five estimators, because this is the depth of the deepest leaves, f and g .

Each factor estimate will determine whether to go *left*, *right*, or *stop*, for each level down the tree. To indicate *stop*, a special vector is used, the identity vector $\mathbf{1}$ (a vector of all ones). By using the appropriate number of these identity vectors, each location in the tree can be thought of as a composition with the same depth (the maximum depth), even if the location is only partially down the tree. For instance, if we consider leaf c in Figure 2, then its position $\mathbf{right} \odot \rho(\mathbf{right}) \odot \rho^2(\mathbf{left})$ is also $\mathbf{right} \odot \rho(\mathbf{right}) \odot \rho^2(\mathbf{left}) \odot \mathbf{1} \odot \mathbf{1}$. This way, we can set up a resonator network for five factors and have it decode locations anywhere in the tree.

We denote each factor estimate as $\hat{\mathbf{x}}^{(0)}$, $\hat{\mathbf{x}}^{(1)}$, $\hat{\mathbf{x}}^{(2)}$, $\hat{\mathbf{x}}^{(3)}$, $\hat{\mathbf{x}}^{(4)}$ and the codebook matrices as \mathbf{X}_0 , \mathbf{X}_1 , \mathbf{X}_2 , \mathbf{X}_3 , \mathbf{X}_4 . Each codebook matrix contains permuted versions of *left* and *right*, and $\mathbf{1}$: $\mathbf{X}_d = [\rho^d(\mathbf{left}), \rho^d(\mathbf{right}), \mathbf{1}]$ where d indicates the depth in the tree. The network is constructed analogous to equation 3.2, but with five factor estimates running in parallel instead of three. For instance, the update equation for the first estimate is

$$\hat{\mathbf{x}}^{(0)}(t+1) = g(\mathbf{X}_0 \mathbf{X}_0^T (\mathbf{s} \odot \hat{\mathbf{x}}^{(1)}(t) \odot \hat{\mathbf{x}}^{(2)}(t) \odot \hat{\mathbf{x}}^{(3)}(t) \odot \hat{\mathbf{x}}^{(4)}(t))). \quad (4.4)$$

The process is demonstrated in Figure 2. The input vector to be factorized, \mathbf{s} , is first formed from the tree data structure and the query. For instance, to find the location of label c , $\mathbf{s} = \mathbf{tree} \odot \mathbf{c}$ is the input to the resonator network. Different leaves in the tree can be found by unbinding the leaf representation from the tree vector and using this result as the input.

We visualize the network dynamics by displaying the similarity of each factor estimate $\hat{\mathbf{x}}^{(d)}(t)$ to the atoms stored in its corresponding codebook \mathbf{X}_d . The evolution of these similarity weights over time is shown as a heat map (see Figure 2, right). The heat maps show that the system initially jumps around chaotically, with the weighting of each estimate changing drastically with each iteration. But then there is a quite sudden transition to a stable equilibrium, where each estimate converges nearly simultaneously,

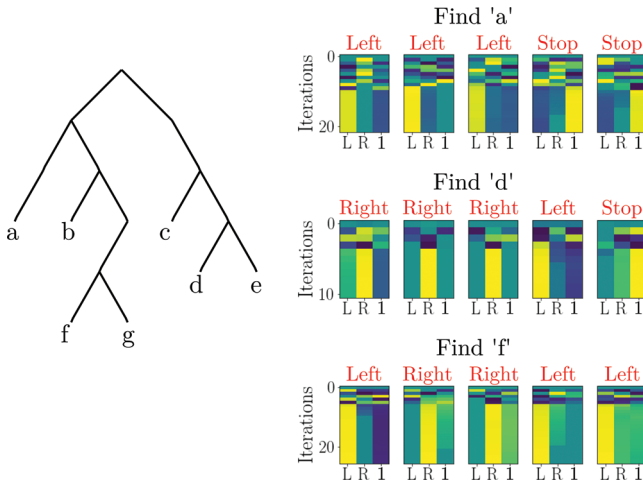


Figure 2: Tree search with a resonator network. The query of the vector **tree** produces an encoding of position that the resonator network can factor. The colored plots indicate the time evolution of $\hat{\mathbf{x}}^{(0)}, \dots, \hat{\mathbf{x}}^{(4)}$ (from left to right), showing the cosine similarity of each estimate to each of the three possible vectors ρ^a (left), ρ^d (right), $\mathbf{1}$. Purple indicates low similarity, and yellow indicates high similarity. Initially the similarity changes significantly until the three estimators find a coherent factorization and quickly converge. Red letters indicate the converged result for each $\hat{\mathbf{x}}^{(0)} \dots \hat{\mathbf{x}}^{(4)}$.

and at this point, the output for each factor is essentially the codebook element with highest weight.

4.2 Visual Scene Analysis as a Factorization Problem. Next, we show how VSAs can encode the compositional structure of a visual scene and how the resonator network can be used to decode the contents of the scene. Consider the scene in Figure 3 containing colored MNIST digits (LeCun, 1998) in different positions. Position in the scene is indexed by vertical and horizontal coordinates, each quantized into three possible values, (top, middle, bottom) and (left, center, right), respectively. Each digit can take on one of seven possible colors (blue, green, cyan, red, pink, yellow, white). The digits are labeled by their semantic class (0, 1, ..., 9), but the exact shape will differ, as the stimuli are sampled from the 50,000 exemplars in the MNIST training set.

Any given scene can have between one and three of these objects, which are allowed to partially occlude one another. We generate symbolic vectors $\mathbf{c}_{\text{blue}}, \mathbf{c}_{\text{green}}, \dots, \mathbf{c}_{\text{white}}$ to encode color; $\mathbf{d}_0, \mathbf{d}_1, \dots, \mathbf{d}_9$ to encode shape; $\mathbf{v}_{\text{top}}, \mathbf{v}_{\text{middle}}, \mathbf{v}_{\text{bottom}}$ to encode vertical position; and




Figure 3: Generating a vector symbolic encoding of a visual scene.

\mathbf{h}_{left} , $\mathbf{h}_{\text{center}}$, $\mathbf{h}_{\text{right}}$ to encode horizontal position, which are stored in respective codebooks, \mathbf{C} , \mathbf{D} , \mathbf{V} , \mathbf{H} .

The example scene (see Figure 3) contains a cyan 7 at position top, left; a pink 3 at position top, right; and a red 8 at position middle, left. While this is a highly simplified type of visual scene, it illustrates the combinatorial challenge of representing and interpreting visual scenes. There are only 23 distinct atomic parameters (10 for digit identity, 7 for color, 3 each for vertical and horizontal position), and yet these combine to describe $10 \times 7 \times 3 \times 3 = 630$ individual objects, and $630 + 630^2 + 630^3 = 250,444,530$ possible scenes with 1, 2, or 3 objects. This number of combinations still does not include the variability among exemplars for each shape, of which there are 50,000 in the MNIST data set.

The VSA approach to represent a scene like this is to form the conjunction of each of the four factors with the binding operation and superposing multiple objects together to form a single high-dimensional vector that constitutes a distributed representation of the entire scene. This encoding is depicted in Figure 3, and as in the previous examples, the encoding provides a flexible data structure such that aspects of the scene can be individually queried. One attractive property of this representation is that its dimensionality does not grow with the number of objects in the scene, nor does it impose any particular ordering on the objects.

To convert a new input image into a structured VSA representation, one challenge is to deal with the variability and correlations between the shapes of different hand-written digits. VSAs are designed for symbolic processing in neural networks. However, when dealing with sensor data streams, one must solve the encoding problem, which is how to map the input data into the symbolic space (Räsänen, 2015; Kleyko, Rahimi, Rachkovskij, Osipov, & Rabaey, 2018). We train a simple feedforward neural network with two fully connected hidden layers to produce the desired VSA encoding of the scene. The feedforward network was trained on a (uniformly) random sample of these scenes, with the MNIST digits chosen from an exclusive training set. A generative model creates the image of the scene from a random sample of factors for each object. From the chosen factors, the VSA representation

of the scene is also generated through binding of VSA vectors for each factor and superposition for each object (see Figure 3). Supervised learning via backpropagation is used to train the network to output the VSA representation of the entire scene from the image pixels as input.

The resonator network can then be used to parse the output of the feed-forward network to identify each object and its properties. The vectors $\hat{\mathbf{c}}(t)$, $\hat{\mathbf{d}}(t)$, $\hat{\mathbf{h}}(t)$ and $\hat{\mathbf{v}}(t)$ denote the guesses for each factor: color, digit, horizontal and vertical-location, respectively. The scene can then be decoded by iterating through the resonator network:

$$\begin{aligned}\hat{\mathbf{c}}(t+1) &= g\left(\mathbf{CC}^\top\left(\mathbf{s} \odot \hat{\mathbf{d}}(t) \odot \hat{\mathbf{v}}(t) \odot \hat{\mathbf{h}}(t)\right)\right), \\ \hat{\mathbf{d}}(t+1) &= g\left(\mathbf{DD}^\top\left(\mathbf{s} \odot \hat{\mathbf{c}}(t) \odot \hat{\mathbf{v}}(t) \odot \hat{\mathbf{h}}(t)\right)\right), \\ \hat{\mathbf{v}}(t+1) &= g\left(\mathbf{VV}^\top\left(\mathbf{s} \odot \hat{\mathbf{d}}(t) \odot \hat{\mathbf{c}}(t) \odot \hat{\mathbf{h}}(t)\right)\right), \\ \hat{\mathbf{h}}(t+1) &= g\left(\mathbf{HH}^\top\left(\mathbf{s} \odot \hat{\mathbf{d}}(t) \odot \hat{\mathbf{c}}(t) \odot \hat{\mathbf{v}}(t)\right)\right).\end{aligned}\tag{4.5}$$

The encoding of visual scenes described superposes a composite vector for each object, each of which individually is a valid solution to the factorization of the scene. When we present the scene vector \mathbf{s} to a resonator network, it automatically hones in on a particular one of these composites, finding its factors. For instance, in Figure 4 the resonator network first identifies the pink 3 at the top right. Once the factorization has been found, this object is then “explained away” by subtracting it from \mathbf{s} . What remains are the other composites, still in superposition. The resonator network is then reset (each resonator is reinitialized to the superposition of all possible codevectors) and presented with the new explained-away scene vector. It will then hone in on one of the remaining objects—in this case, the red 8. This sequence may be repeated until all the objects have been decoded. This technique is similar to what is known as deflation in the context of tensor decomposition methods (da Silva, Comon, & de Almeida, 2015).

After training on 100,000 images, we used the network to produce symbolic vectors for a held-out test set of 10,000 images. The vector dimensionality N is a free parameter, which we chose to be 500. If the exact ground-truth vector is provided to a resonator network, it will infer the factors with 100% accuracy provided N is large enough, a fact we establish in the companion article in this issue. For this small, visual scene example, it turns out $N = 500$ more than suffices for the number of possible factorizations to be searched. Note that $N = 500$ is lower than the total number of combinations of all the factors, which is 630.

The encoder network generates VSA scene vectors that are close to the ground-truth encoding, but there is some error. The error gets larger with more digits in the scene, perhaps partially due to occlusion of the digits.

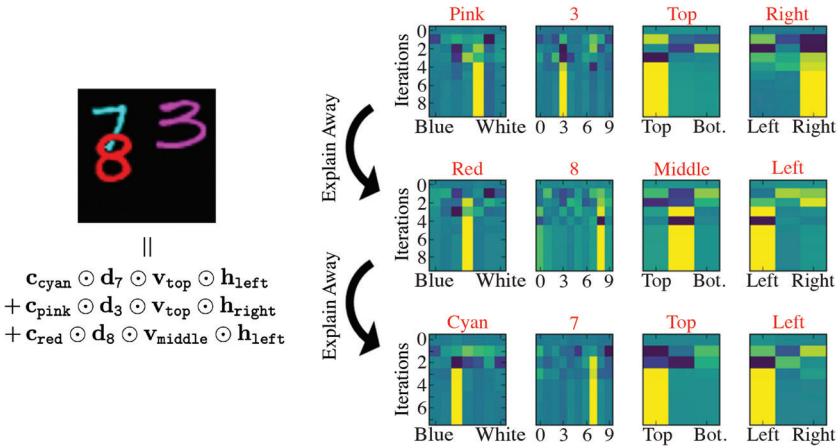


Figure 4: Scene vector s is fed into a resonator network that decodes each object in the scene. The model hones in on one object at a time, which is then explained away by subtracting the resonator network’s converged state from the scene vector. The network is reset and provided with this new input vector. It then converges to another solution, which describes a different object in the scene.

Figure 5 shows that the resonator network can tolerate significant error in the scene vector produced by the feedforward encoding network, correcting for ambiguity not resolved in the encoding step.

5 Discussion

A major quest for modern artificial intelligence is to build computational models that combine the abilities of neural networks with the abilities of rule-based reasoning. Vector Symbolic Architectures, a family of connectionist models, enable the formation of distributed representations of data structures, structured computation on these representations, and has provided valuable conceptual insights for cognition and computation. However, so far, VSA models have not been able to solve challenging artificial intelligence problems in real-world domains due to the combinatorial factorization problem that arises when processing complex, hierarchical data structures. Our contribution here has been to provide an efficient solution to the factorization problem, the resonator network, which we show in the companion article in this issue (Kent et al., 2020) vastly outperforms standard optimization methods.

The two applications we showed here—parsing a tree-like data structure and decomposing a visual scene—are intended as illustrative examples to show how factorization of multiterm products arises in querying a VSA data structure, and they show how to design resonator networks to solve

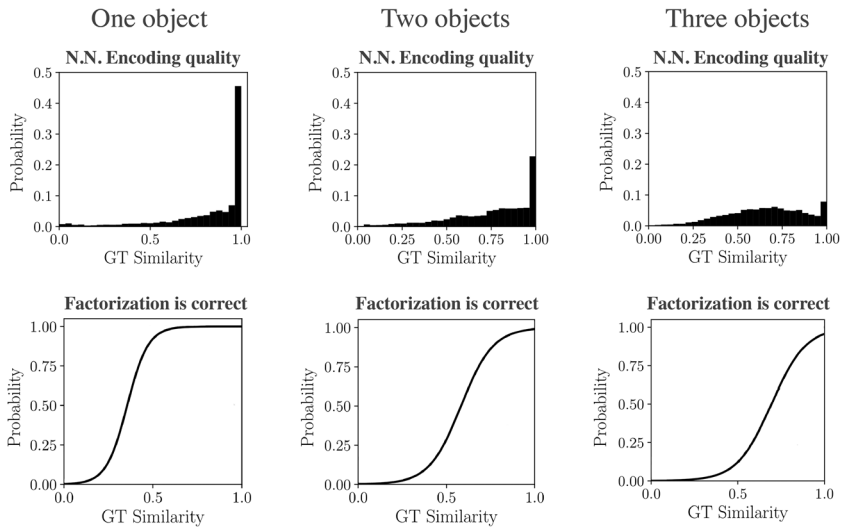


Figure 5: Resonator networks correct encoding errors. Visual scenes with one, two, and three objects are separated into separate columns. (Top) Encoding quality in terms of cosine similarity between the feedforward network output and the ground-truth scene vector, across the test set. We define correct factorization as the case where the resonator network correctly infers all the factors of all objects. (Bottom) The empirical probability of a correct factorization as a function of similarity to the ground-truth scene vector. Lines are logistic function fits to the data.

such problems. Having a solution to the factorization problem now makes it possible to apply VSAs to myriad problems in computational neuroscience, cognitive science, and artificial intelligence, from visual scene analysis to natural language understanding and analogical reasoning.

5.1 Implications for Neuroscience. The ability to solve factorization problems is fundamental to both perception and cognition. In vision, for example, the signal measured by a photoreceptor contains a combination of illumination, surface reflectance, surface orientation, and atmospheric properties that essentially need to be “demultiplied” by the visual system in order to recover a representation of the underlying causes in a scene (Barrow & Tenenbaum, 1978; Adelson & Pentland, 1996; Barron & Malik, 2014). The problem of separating form and motion may also be posed as a factorization problem (Cadieu & Olshausen, 2012; Memisevic & Hinton, 2010; Anderson et al., 2020). In the domain of language, it has been argued that a factorization of sentence structure into “roles” and “fillers” is required for robust and flexible processing (Smolensky, 1990; Jackendoff,

2002). Many cognitive tasks, such as analogical reasoning, also require a form of factorization (Hummel & Holyoak, 1997; Kanerva, 1998; Plate, 2000a; Gayler & Levy, 2009). However, to date, it has been unclear how these factorization problems could be represented and solved efficiently by neural circuits in the brain. VSAs and resonator networks are a potential neural solution to these problems, and indeed developing more neurobiologically plausible models along these lines is a goal of ongoing work.

In the context of neuroscience and psychology, binding is widely theorized to be an important process by which the brain properly associates features belonging to the same physical object. However, how the brain may accomplish this is a hotly debated subject. Various solutions to this problem, also known as the neural binding problem, have been proposed based on attentional mechanisms or neural synchrony (Treisman & Gelade, 1980; von der Malsburg, 1999; Wolfe & Cave, 1999). Note that in these proposals, the binding information required to properly describe sets of compound objects has to be added to the individual feature representations, thus increasing the dimension for representing a compound object (or expanding the representation in time).

VSAs provide a general solution to the binding problem that early visual stages could employ. By using the VSA operations to represent and form data structures, the binding of features is easily expressed. Further, the dimension of the compound representation is not increased. The main computational challenge then becomes the factorization of VSA data structures formed in early sensory pathways, for which the resonator network provides a neurally plausible solution. Interestingly, Feldman's (2013) earlier discussion of the binding problem has already pointed out that the more fundamental problem of sensory processing is actually one of *unbinding*. Feldman argued that the raw sensory signals themselves can be thought of as being composites, containing multiple attributes that require factorization, such as in the examples described here.

In terms of modeling computation in biological neural circuits, resonator networks are clearly an abstraction. In particular, the implementation of VSAs presented here assumes that information is encoded by dense bipolar vectors (each element is nonzero), and the binding operation is performed by element-wise multiplication of vectors. At first glance, these types of representations and operations may not seem very biologically plausible. However, other variants of VSAs that utilize sparse, rather than dense, representations may help to reconcile this disconnect (Rachkovskij & Kussul, 2001; Laiho, Poikonen, Kanerva, & Lehtonen, 2015). Recently, we have shown that compound objects can be efficiently represented by sparse vectors with the same dimension as the atomic representations (Frady, Kleyko, & Sommer, 2020). The binding operation in this context relies on sigma-pi type operations (Mel & Koch, 1990; Plate, 2000b) that are potentially compatible with active nonlinearities found in dendritic trees. Complex-valued

variations of VSAs (Plate, 2003) can also be linked to spike-timing codes (Frady & Sommer, 2019), which could further increase links to biology.

5.2 Implications for Machine Learning. In conventional deep learning approaches, given enough labeled data, a multilayer network can be trained end-to-end without worrying about understanding or parsing the representations formed by the intermediate layers. Users typically consider the interior of a deep network as a black box. However, this conceptual convenience becomes a disadvantage when it comes to improving the deficiencies of deep learning methods: susceptibility to adversarial attacks, the need for large amounts of labeled data, and a lack of generalization to novel situations. Moreover, while most machine-learning algorithms are focused on problems of pattern matching or learning a mapping from inputs to outputs, most problems in perception and cognitive reasoning require more than just pattern matching; they also the ability to form and manipulate data structures.

VSAs offer a transparent approach to forming distributed representations of potentially complex data structures that may be flexibly recombined to deal with novel situations. For any desired computation, the relevant elements in the data structure can be exposed, or decoded, and combined with other information to calculate a result. Here we have shown how these data structures can be formed and manipulated to solve challenging computational problems such as tree search or visual scene analysis. The key to solving these problems relies on the ability to factorize high-dimensional vectors, which can now be done by resonator networks. Given that the problem of factorization arises in many other machine-learning settings, such as simultaneous inference of multiple variables, inverse graphics, and language, it seems likely that resonator networks could provide an efficient solution in these domains as well.

One can potentially combine VSAs with deep learning to get the best of both worlds. An example of this may be seen in our solution to parsing a visual scene (see section 4.2). Rather than training a network to simply map images to class labels, our approach trains the network to map the image to a symbolic description that captures the compositional structure of a scene—that is, multiple objects combined with their properties—which can be used by downstream processes to reason about the scene. Importantly, because multiple object-property bindings can be superposed in the same space, the VSA encoding can handle the very large combinatoric space of possible scenes (in this case, 250 million) with a single vector of fixed dimensionality (500). The VSA representation does have a limited capacity and will begin to break down for more than a few objects. However it is worth noting that human working memory has similar limitations (Miller, 1956).

While there are undoubtedly alternative deep learning approaches for performing analysis of simple scenes, our goal here was to show how

analysis of visual scenes could be approached by expressing the problem as a problem of factorization. Incorporating factorization into problems like scene analysis may enable reasoning in much more complex spaces, as such a system can utilize factorization to handle a very large combinatoric space. However, the simple hybrid approach presented here still has some shortcomings, such as requiring a large amount of training data to learn the encoding.

We believe that multilayer neural networks could be improved profoundly by enabling all layers to explicitly represent, learn, and factorize data structures. Some recent model innovations follow this direction, particularly the “transformer” neural network architecture, which encodes key-value pairs for modeling language and other types of data (Vaswani et al., 2017; Devlin, Chang, Lee, & Toutanova, 2018). Other model proposals enable the encoding of multiplicative relationships between features using the tensor product (Nickel, Tresp, & Kriegel, 2011; Socher, Chen, Manning, & Ng, 2013). VSAs could enable these models to represent and manipulate increasingly complex data structures, but this requires solving factorization problems. Resonator networks could thus serve as a critical component for building trainable neural networks that form, query, and manipulate large hierarchical data structures.

Acknowledgments

We thank members of the Redwood Center for Theoretical Neuroscience for helpful discussions, in particular Pentti Kanerva, whose work on Vector Symbolic Architectures originally motivated this project. This work was generously supported by NIH grant 1R01EB026955-01, NSF grants IIS1718991 and DGE1752814, the Intel Neuromorphic Research Community, Berkeley Deep-Drive, the Semiconductor Research Corporation and NSF under E2CDA-NRI, DARPA’s Virtual Intelligence Processing program, and AFOSR FA9550-19-1-0241.

References

- Adelson, E. H., & Pentland, A. P. (1996). The perception of shading and reflectance. In D. Knill & W. Richards (Eds.), *Perception as Bayesian inference* (pp. 409–423). New York: Cambridge University Press.
- Anderson, A. G., Ratnam, K., Roorda, A., & Olshausen, B. (2020). High acuity vision from retinal image motion. *Journal of Vision*, 20(7), 34.
- Barron, J. T., & Malik, J. (2014). Shape, illumination, and reflectance from shading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(8), 1670–1687.
- Barrow, H., & Tenenbaum, J. (1978). Recovering intrinsic scene characteristics from images. *Computer Vision Systems*, 2, 3–26.
- Cadieu, C. F., & Olshausen, B. A. (2012). Learning intermediate-level representations of form and motion from natural movies. *Neural Computation*, 24(4), 827–866.

- Cox, G. E., Kachergis, G., Recchia, G., & Jones, M. N. (2011). Toward a scalable holographic word-form representation. *Behavior Research Methods*, 43(3), 602–615.
- da Silva, A. P., Comon, P., & de Almeida, A. L. (2015). An iterative deflation algorithm for exact CP tensor decomposition. In *Proceedings of the 2015 IEEE International Conference on Acoustics, Speech and Signal Processing* (pp. 3961–3965). Piscataway, NJ: IEEE.
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). *Bert: Pre-training of deep bidirectional transformers for language understanding*. arXiv:1810.04805.
- Feldman, J. (2013). The neural binding problem(s). *Cognitive Neurodynamics*, 7, 1–11.
- Fodor, J. A. (1975). *The language of thought*. New York: Crowell.
- Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1–2), 3–71.
- Frady, E. P., Kleyko, D., & Sommer, F. T. (2018). A theory of sequence indexing and working memory in recurrent neural networks. *Neural Computation*, 30, 1449–1513.
- Frady, E., Kleyko, D., & Sommer, F. (2020). *Variable binding for sparse distributed representations: Theory and applications*. arXiv:209.06734.
- Frady, E. P., & Sommer, F. T. (2019). Robust computation with rhythmic spike patterns. In *Proceedings of the National Academy of Sciences*, 116(36), 18050–18059.
- Gayler, R. W. (1998). Multiplicative binding, representation operators and analogy [workshop poster]. In K. Holyoak, D. Gentner, & B. Kokinov (Eds.), *Advances in analogy research: Integration of theory and data from the cognitive, computational, and neural sciences*. Berlin: Springer.
- Gayler, R. (2003). Vector symbolic architectures answer Jackendoff's challenges for cognitive neuroscience. In *Proceedings of the ICCS/ASCS International Conference on Cognitive Science*. Amsterdam: Elsevier Procedia.
- Gayler, R. W., & Levy, S. D. (2009). A distributed basis for analogical mapping. In *New Frontiers in Analogy Research: Proceedings of the Second International Analogy Conference-Analogy*, vol. 9 (pp. 165–174). NBU Press.
- Hinton, G. E. (1990). Mapping part-whole hierarchies into connectionist networks. *Artificial Intelligence*, 46(1–2), 47–75.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. In *Proceedings of the National Academy of Sciences*, 79(8), 2554–2558.
- Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, 104(3), 427.
- Jackendoff, R. (2002). *Foundations of language: Brain, meaning, grammar, evolution*. Oxford: Oxford University Press.
- Joshi, A., Halseth, J. T., & Kanerva, P. (2016). Language geometry using random indexing. In *Proceedings of the International Symposium on Quantum Interaction* (pp. 265–274). Berlin: Springer.
- Kanerva, P. (1996). Binary spatter-coding of ordered k-tuples. In *Proceedings of the International Conference on Artificial Neural Networks* (pp. 869–873). Berlin: Springer.
- Kanerva, P. (1997). Fully distributed representation. In *Proceedings of the 1997 Real World Computing Symposium* (pp. 358–365). Real World Computing Partnership.

- Kanerva, P. (1998). Large patterns make great symbols: An example of learning from example. In *Proceedings of the International Workshop on Hybrid Neural Systems* (pp. 194–203). Berlin: Springer.
- Kanerva, P. (2009). Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors. *Cognitive Computation*, 1(2), 139–159.
- Kent, S. J., Frady, E. P., Sommer, F. T., & Olshausen, B. A. (2020). Resonator networks, 2: Factorization performance and capacity compared to optimization-based methods. *Neural Computation*, 32(12), 2332–2388.
- Kleyko, D., Rahimi, A., Rachkovskij, D. A., Osipov, E., & Rabaey, J. M. (2018). Classification and recall with binary hyperdimensional computing: Tradeoffs in choice of density and mapping characteristics. *IEEE Transactions on Neural Networks and Learning Systems*, 29(12), 5880–5898.
- Laiho, M., Poikonen, J. H., Kanerva, P., & Lehtonen, E. (2015). High-dimensional computing with sparse vectors. In *Proceedings of the 2015 IEEE Biomedical Circuits and Systems Conference* (pp. 1–4). Piscataway, NJ: IEEE.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40.
- LeCun, Y. (1998). *The MNIST database of handwritten digits*. <http://yann.lecun.com/exdb/mnist>
- Ledoux, M. (2001). *The concentration of measure phenomenon*. Providence, RI: American Mathematical Society.
- McClelland, J. L., Rumelhart, D. E., & PDP Research Group. (1986). *Parallel distributed processing: Explorations in the Microstructure of Cognition* 1. Cambridge, MA: MIT Press.
- Mel, B. W., & Koch, C. (1990). Sigma-Pi learning: On radial basis functions and cortical associative learning. In D. S. Touretzky (Ed.), *Advances in neural information processing systems*, 2 (pp. 474–481). San Mateo, CA: Morgan Kaufmann.
- Memisevic, R., & Hinton, G. E. (2010). Learning to represent spatial transformations with factored higher-order Boltzmann machines. *Neural Computation*, 22(6), 1473–1492.
- Miller, G. A. (1956). The magic number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, 91–97.
- Nickel, M., Tresp, V., & Kriegel, H.-P. (2011). A three-way model for collective learning on multi-relational data. In *Proceedings of the 28th International Conference on Machine Learning*.
- Plate, T. A. (1991). Holographic reduced representations: Convolution algebra for compositional distributed representations. In *Proceedings of the International Joint Conference on Artificial Intelligence* (pp. 30–35). San Mateo, CA: Morgan Kaufmann.
- Plate, T. A. (1995). Holographic reduced representations. *IEEE Transactions on Neural Networks*, 6(3), 623–641.
- Plate, T. A. (2000a). Analogy retrieval and processing with distributed vector representations. *Expert Systems*, 17(1), 29–40.
- Plate, T. A. (2000b). Randomly connected sigma-pi neurons can form associator networks. *Network: Computation in Neural Systems*, 11(4), 321–332.

- Plate, T. A. (2003). *Holographic reduced representation: Distributed representation of cognitive structure*. Stanford, CA: CSLI Publications.
- Rachkovskij, D. A., & Kussul, E. M. (2001). Binding and normalization of binary sparse distributed representations by context-dependent thinning. *Neural Computation*, 13(2), 411–452.
- Räsänen, O. J. (2015). Generating hyperdimensional distributed representations from continuous-valued multivariate sensory input. In *Proceedings of the 37th Annual Meeting of the Cognitive Science Society*. Red Hook, NY: Curran.
- Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence*, 46(1–2), 159–216.
- Socher, R., Chen, D., Manning, C. D., & Ng, A. (2013). Reasoning with neural tensor networks for knowledge base completion. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems*, 26 (pp. 926–934). Red Hook, NY: Curran.
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12(1), 97–136.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., . . . Polosukhin, I. (2017). Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Advances in neural information processing systems*, 30 (pp. 5998–6008). Red Hook, NY: Curran.
- von der Malsburg, C. (1999). The what and why of binding: The modeler's perspective. *Neuron*, 24(1), 95–104.
- Wolfe, J. M., & Cave, K. R. (1999). The psychophysical evidence for a binding problem in human vision. *Neuron*, 24(1), 11–17.

Received March 2, 2020; accepted July 6, 2020.