A common framewor, for distri,uted representation scfor compositional Tony A. School of Mat

representation for a structure: a reduced representation is a

A framework for distributed representation In the last few years people have proposed a number of schemes for di

from a very large pool or have complex In local representatio

A framework for distributed representation a particular entity. In distributed representations entities are represen

- Distributed representations offer an efficient use of representational resources (provided that one only wants to represent a small number of entities at any one time, and that some degree of overlap is permissible), because the number of possible patterns over a collection of units is far greater than the number of units in the collection.
- Distributed representations are continuous, or at least fine-grained. This allows nuances of meaning to be expressed by small changes in patterns, and supports learning.
- Distributed representations can be processed in fine-grained parallel networks of simple processing units (i.e., neural networks). Furthermore it is possible to learn or modify patterns by making small changes to them with gradient-based techniques for learning in neural networks.
- Distributed representations can be noise tolerant, as patterns are still recognizable when slightly corrupted.

where A n

- Patterns can be superimposed, but still individually recognizable. As more patterns are superimposed, it becomes more difficult to determine which patterns are present. As a result, connectionist memories which use distributed representations tend to have soft capacity limits.
- Using distributed representations it is possible to capture a rich similarity structure over entities, in which surface similarity of patterns is sensitive to both structure and non-structural features of entities (see Gentner and Markman [1993] and Plate [1994a] for further

3A common framework

All the abovementioned schemes for representing composit

etc. are the individual values, and n is number of elements in

−−− it randomizes unstructured similarity: **A** −−− **B** is not similar to **A** or **B**;

−−− it preserves structured similarity: **A** −−− **B** is similar to **A**′ −−− **B**′ to the extent that **A**′ is similar to **A** and **B**′ is similar to **B**; and

−−− an inverse (decoding) operation exists, so that the pattern **B** can be reconstructed given **A** −−− **B** and **A** (or vice-vBinding operation diA23er more than superposition operations. The tw

2. the choice of superposition operation;

“John ate the A24sh”, can be represented as: **eat agent** −−− **john** + **eat object** −−− **the Afifish**. A al-terna

–

1. the nature of the distributed representation: the domain of activation values (binary, real, or complex), and the statistical distribution of pattern values; whether representations are fully or partially distributed;

3. the choice of binding operation (which determines whether or not bindings are the same type of representational objects as patterns);
4. how the binding operation is used to represent predicate structure (role-A24ller binding is a common method, but there are others); and

5. the use of other operations and techniques, such as non-linearities, clean-up memory, normalization, and **Why combinatorial features are practical with distributed representations**

–

–

–

–

–

${}_iB_j$.**4.1 Holographi**

$z_i = {}^{n}$

Scheme	Patterns	Superposition	Binding	Structure	Other operations
HRRs	real (Gaussian)	addition	circular convolution		

$$\begin{aligned}\mathbf{P}_{objects} &= \langle \mathbf{jane} + \mathbf{spot} \rangle \\ \mathbf{P}_{cause} &= \langle \mathbf{cause} + \mathbf{P}_{objects} + \mathbf{P}_{bite} + \mathbf{P}_{flee} \\ &\quad + \mathbf{cause}_{antc} \circledast \mathbf{P}_{bite} + \mathbf{cause}_{cnsq} \circledast \mathbf{P}_{flee}\end{aligned}$$

The unbound representations for the ob

$_{objects}$) and lower-level HRRs ($\mathbf{P}_{bite} + \mathbf{P}_{flee}$)are included in $\mathbf{P}$

$_{cause}$ in order to make the representation for a proposition have some simi-larity to those objec

$_{fr}$):

$$\begin{aligned}\mathbf{jane} &= \langle \mathbf{person} + \mathbf{id}_{jane} \rangle \\ \mathbf{fred} &= \langle \mathbf{person} + \mathbf{id}_{fred} \rangle\end{aligned}$$

The mixing proportions of the various components

of the role (or A24ller)pattern. For exar

$\circledast$ **spot** can be found by convolving it with**bite** multiple bindings are superimposed, as in **P**

$_{bite}$, the other bindings do not interfere withthe decoding other strategies used. The only components which count towards the size

capacity to be interesting and potentially very useful.

4.4 CADAM, TODAM, TODAM2 and

Liepa's [1977] CAI

[5] Murdock's TODAM[6] [Murdock 1982] and TODAM2 [Murdoc1993], and Metcalf-E

[7] [1982, 1985] use other forms of convolution to items in models of hum

$^{\mathrm{T}}$. T has n^2 elements, and $T_{ij} = x_i y_j$. Both lower-order higher-order tensors exi

$_{ijk} = x_i y_j z_k$ and T has $.3$ elements. Higher-order tensor products are constructed in asimilar fashion – a

$_k$elements. Tensor

CADAM

A framework for distributed representation [1990], Dolan and Smolensky [1989], and Dolan [1989] use role-A24ller

nesting. Tensor products bindings are one extreme of using the pairwise products of pat-tern elements: ever

Care must be exercised to take the appropriate inner product, which depends upon the position of unknown component in th

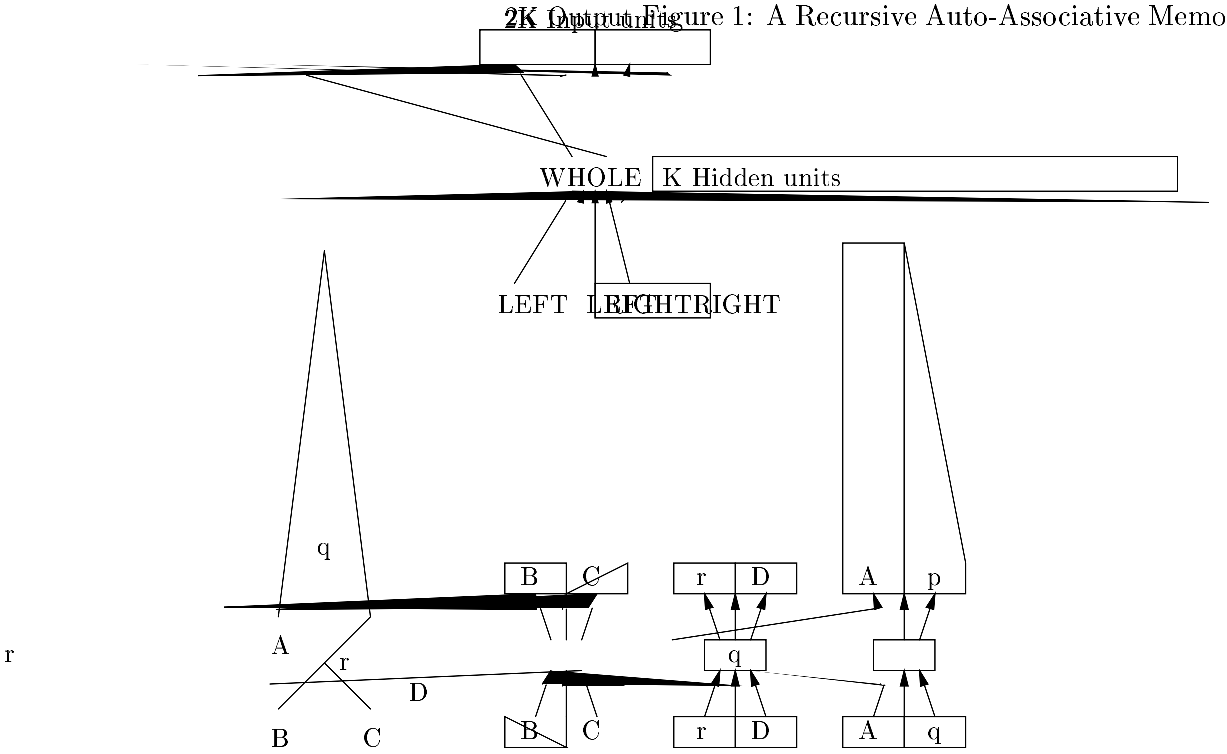

Figure 1: A Recursive Auto-Associative Memory (RAAM) for bi

Figure 2: A simple tree and the auto-associations that encode it in a similar if any of the role-A24ller p

half), and a decoding net (the top half). The encoding net converts a full representation toa reduced representa representation for the tree (B C) is r, the reduced representation for (

Both encoding and decoding are recursive operations. The encoding procedure knows, fromthe structure of th

Smolensky [personal communication] observed that RAAMs can be viewed as implementingrole-filler binding

binding operation is the sum of certain elements of the tensor produ

vector products of roles and fillers. The weights from the hidden to output units mustimplement the in

——— Bindings are “first class” objects? If **A** ——— **B** is the same type of pattern as **A** and **B** (or at least the same as one of them), then it is straightforward to store nested compositional structure using role-filler bindings. This is the case for HRRs, binary spatter codes, and RAAMs. Chunking is simple to apply when bindings are the same size as the original patterns. However, with tensor products, bindings grow exponentially in size with the depth of structure.

——— Is structure represented by role/filler bindings (HRRs, binary spatter codes, tensor products, RAAMs) or by argument-product bindings (STAR)? This has important consequences for the properties of the model.

——— Status of role patterns. In RAAMs, role patterns cannot vary (once the RAAM net has been trained), whereas in the other role-filler schemes, role patterns can be different for different predicates. Furthermore, in these other schemes, role patterns can be made similar or different. For example, we might want the role patterns for the agent roles of “bite” and “chew” to be quite similar, and the agent roles of “eat” and “sell” to be quite different. This will have consequences for the structure-sensitive operations.

——— Algebraic properties of superposition and binding operations. Ordinary and wrapped convolution are commutative (**A** ——— **B** ——— **B** ——— **A**) and associative ((**A** ——— **B**) ——— **C** ——— **A**

C)), whereas the tensor product is only associative. This has consequences for how structures are decoded: convolution structures can be decoded in any order, but tensor product structures must be decoded in a particular order. The non-commutativity of the tensor product is essential to the STAR model: if STAR were implemented with a commutative binding operation the order of arguments would be lost. Non-commutative and non-associative versions of convolution do exists, see Plate **5.1 Estimating struct**

represenGentner and Markman, [Gentner and Markman 1993] expressed skepticism regarding possib

A framework for distributed representation requires the matching of descriptions at the same level of abstract

A framework for distributed representation without dereferencing any pointers. This is what the “reduced” in “

The networks learned to work effectively with much lower dimensional patterns than areneeded in a syste

Similar learning techniques can also be used with other methods. For example, Plate constructed recurrent netw

sums of these elements can also be an effective binding operation. The properties of such a

Chalmers, D. J. (1990). Syntactic transformations on distributed representations. *Connection Science 2*(1 & 2), 53–62.

Dolan, C. P. (1989). Tensor manipulation networks: Connectionist and symbolic approaches to comprehension, learning, and planning. Computer Science Department, AI Lab. Technical Report UCLA-AI-89-06, UCLA.

Dolan, C. P. and P. Smolensky (1989). Tensor product production system: a modular architecture and representation. *Connection Science 1(1)*, 53–68.

Feldman, J. A. (1986). Neural representation of conceptual knowledge. Technical report TR189, Department of Computer Science, University of Rochester, Rochester NY.

Gentner, D. and A. B. Markman (1993). Analogy – Watershed or Waterloo? Structural alignment and the development of connectionist models of analogy. In C. L. Giles, S. J. Hanson, and J. D. Cowan (Eds.), *Advances in Neural Information Processing Systems 5 (NIPS*92)*, San Mateo, CA, pp. 855–862. Morgan Kaufmann.

Halford, G. S., W. H. Wilson, J. Guo, R. W. Gayler, J. Wiles, and J. E. M. Stewart (1994). Connectionist implications for processing capacity limitations in analogies. In K. J. Holyoak and J. Barnden (Eds.), *Analogical Connections*, Volume 2 of *Advances in Connectionist and Neural Computation Theory*. Norwood, NJ: Ablex.

Hinton, G. E. (1990). Mapping part-whole heirarchies into connectionist networks. *Artificial Intelligence 46*(1-2), 47–76.

Hinton, G. E., J. L. McClelland, and D. E. Rumelhart (1986). Distributed representations. In D. E. Rumelhart, J. L. McClelland, and the PDP research group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition.*, Volume 1, pp. 77–109. Cambridge, MA: MIT Press.

Hummel, J. E. and I. Biederman (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review 99*(3), 480–517.

Kanerva, P. (1996). Binary spatter-coding of ordered k-tuples. In C. von der Malsburg, W. von Seelen, J. Vorbruggen, and B. Sendhoff (Eds.), *Artificial Neural Networks–ICANN Proceedings*, Volume 1112 of *Lecture Notes in Computer Science*, Berlin, pp. 869–873. Springer.

Legendre, G., Y. Miyata, and P. Smolensky (1991). Distributed recursive structure processing. In D. Touretzky and R. Lippman (Eds.), *Advances in Neural Information Processing Systems 3*, San Mateo, CA, pp. 591–597. Morgan Kaufmann.

Liepa, P. (1977). Models of content addressable distributed associative memory. Unpublished manuscript.

Metcalfe Eich, J. (1982). A composite holographic associative recall model. *Psychological Review 89*, 627–661.

Metcalfe Eich, J. (1985). Levels of processing, encoding specificity, elaboration, and charm. *Psychological Review 92*, 1–38.

Murdock, B. B. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review 89*(6), 316–338.

Murdock, B. B. (1983). A distributed memory model for serial-order information. *Psychological Review 90*(4), 316–338.

Murdock, B. B. (1987). Serial-order effects in a distributed-memory model. In D. S. Gorfein and R. R. Hoffman (Eds.), *MEMORY AND LEARNING: The Ebbinghaus Centennial Conference*, pp. 277–310. Lawrence Erlbaum Associates.

Ph. D. thesis, Department of Computer Science, University of Toronto. Available at

Murdock, B. B. (1993). TODAM2: A model for the storage and retrieval of item, associative, and serial-order information. *Psychological Review 100*(2), 183–203.

Niklasson, L. and T. van Gelder (1994). Can connectionist models exhibit non-classical structure sensitivity? In *Proceedings of the Sixteenth Annual Conference of The Cognitive Science Society.*

Plate, T. A. (1991). Holographic Reduced Representations: Convolution algebra for compositional distributed representations. In J. Mylopoulos and R. Reiter (Eds.), *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, San Mateo, CA, pp. 30–35. Morgan

Plate, T. A. (1994a). *Distributed Representations and Nested Compositional Structure.*

http://www.cs.utoronto.ca/~tap.

Plate, T. A. (1994b). Estimating structural similarity by vector dot products of holographic reduced representations. In J. D. Cowan, G. Tesauro, and J. Alspector (Eds.), *Advances in Neural Information Processing Systems 6 (NIPS*93)*, San Mateo, CA, pp. 1109–1116. Morgan Kaufmann.

Plate, T. A. (1997). Structure matching and transformation with distributed representations. In R. Sun and F. Alexandre (Eds.), *Connectionist Symbolic Integration.* Lawrence Erlbaum Associates.

Pollack, J. B. (1990). Recursive distributed representations. *Artificial Intelligence 46*(1-2), 77–105.

Press, W. H., B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling (1992). *Numerical Recipes in C* (Second ed.). Cambridge: Cambridge University Press.

Shastri, L. and V. Ajjanagadde (1993). From simple associations to systematic reasoning: A connectionist representation of rules, variables and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences 16*(3), 417–494.

Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence 46*(1-2), 159–216.

Sperduti, A. (1994a). Encoding of labeled graphs by labeling raam. In J. D. Cowan, G. Tesauro, and J. Alspector (Eds.), *Advances in Neural Information Processing Systems 6*, San Mateo, CA, pp. 1125–1132. Morgan Kaufmann.

Sperduti, A. (1994b). Labeling raam. *Connection Science 6*(4), 429–459.

Sperduti, A. and A. Starita (1995). Supervised neural networks for the classification of structures. Technical Report TR-16/95, University of Pisa, Dipartimento di Informatica.

Sperduti, A., A. Starita, and C. Goller (1997). Distributed representations for terms in hybrid reasoning systems. In R. Sun and F. Alexandre (Eds.), *Connectionist Symbolic Integration.* Lawrence Erlbaum Associates.

Touretzky, D. S.