Holographic Reduced Representations

Tony Plate
Department of Computer Science, University of Toronto
Toronto, Ontario, Canada, M5S 1A4
tap@ai.utoronto.ca

Abstract

Associative memories are conventionally used to represent data with very simple structure: sets of pairs of vectors. This paper describes a method for representing more complex compositional structure in distributed representations. The method uses circular convolution to associate items, which are represented by vectors. Arbitrary variable bindings, short sequences of various lengths, simple frame-like structures, and reduced representations can be represented in a fixed width vector. These representations are items in their own right, and can be used in constructing compositional structures. The noisy reconstructions extracted from convolution memories can be cleaned up by using a separate associative memory that has good reconstructive properties.

I Introduction

Distributed representations [13] are attractive for a number of reasons. They offer the possibility of representing concepts in a continuous space, they degrade gracefully with noise, and they can be processed in a parallel network of simple processing elements. However, the problem of representing compositional structure[1] in distributed representations has been for some time a prominent concern of both proponents and critics of connectionism [9, 32, 12].

Most work on neural-network style associative memories has focussed on either auto-associative or hetero-associative memories. Auto-associative memories, e.g., Hopfield networks [14], store an unordered set of items. They can be used to recall item given a distorted version of one. Hetero-associative memories, e.g., holographic memories and matrix memories [37, 8, 22, 5, 38], store a set of pairs of items. One item of a pair can be recalled using the other as a cue. Matrix style memories are the more popular class, owing to superior storage capacity and fewer constraints on vectors to be stored.

For artificial intelligence tasks such as language processing and reasoning the need arises to represent more complex data structures such as sequences and trees. It is difficult to represent sequences or trees in distributed representations using associations of pairs (or even n-tuples) of items and retain the benefits of distributed representations. The problem with representing compositional structure in most associative memories is that items and associations are represented in different spaces. For example, in a Hopfield memory (a matrix style memory) items are represented on unit activations (a vector) and associations are represented on connections weights (a matrix). This makes it difficult to represent relationships with recursive structure in which an association of items may be the subject of another association.

Hinton [12] discusses this problem and proposes a framework in which "reduced descriptions" are used to represent parts and objects in a part-whole hierarchy (a frame-like representation). This framework requires that a number of vectors, each a part and together forming a whole, be compressed (reduced) into a single vector of the same dimension as the original vectors. This reduced vector can in turn be used as a part in the representation of some greater whole. The reduction must be reversible so that one can move in both directions in a part-whole hierarchy, i.e., reduce a set of vectors (a whole) to a single vector (a potential part), and expand a single vector (a part) to a set of vectors (a whole). In this way, compositional structure is represented. An essential aspect of reduced descriptions is that they should be systematically related to their components, so that information about the components can be gleaned without expansion. It is this aspect that distinguishes reduced descriptions from arbitrary pointers. Unfortunately, Hinton does not suggest any concrete way of performing the reduction and expansion mappings.

Some researchers have built models or designed frameworks in which some compositional structure is present in distributed representations. For some examples see the papers of Touretzky [33], Pollack [27], or Smolensky [32].

In this paper I propose a new method for representing compositional structure in distributed representations. Circular convolutions are used to construct associations of vectors. The representation of an association is a vector of the same dimensionality as the vectors which are associated. This allows the construction of representa-

[1] I.e., recursive, or tree-like structure.

IEEE Transactions on Neural Networks 6:3:623-641 1995