Learning of Compositional Hierarchies By Data-Driven Chunking

Karl Pfleger
Computer Science Department
Stanford University
Stanford, CA 94305
kpfleger@cs.stanford.edu

Compositional hierarchies (CHs), layered structures of *part-of* relationships, underlie many forms of data, and representations involving these structures lie at the heart of much of AI. Despite this importance, methods for learning CHs from data are scarce. We present an unsupervised technique for learning CHs by an on-line, bottom-up *chunking* process. At any point, the induced structure can make predictions about new data.

Currently we work with with 1-dimensional, discrete-position data (sequential data but not necessarily directional), where each position is occupied by a symbol from a discrete, finite alphabet, and where the data is potentially unbounded in both directions, not organized into a set of strings with definite beginnings and ends. Unlike some chunking systems, our chunking is purely data-driven in the sense that only the data, and its underlying statistical regularities, are used to construct the hierarchy. Abstractly, we proceed by repeatedly identifying frequently occurring substructures and forming new nodes for them in the hierarchy.

We have developed two systems, both implemented as symmetric, recurrent neural networks. Both can be viewed as follow-ups to the successful interactive activation model (IAM) (McClelland & Rumelhart 1981), extending it in a number of ways, primarily to incorporate learning. The first system and more details on CH learning in general are described in (Pfleger 1998).

The second system uses a Boltzmann machine, extended to handle categorical values and weight sharing. As with IAM, the network encodes a CH directly using a localist representation. Weight sharing and "hardware" duplication (unrolling) are used to model atomic symbols and chunks at different positions. Atomic symbols (the data) are visible variables and the chunks are hidden variables. Chunking is accomplished as an on-line structure modification rule. Specifically, new chunks are created when the weights between existing atoms or chunks show strong positive co-occurrence. Somewhat like a reverse-decimation, the original direct weight is removed and a new node representing the aggregation is created with links to its two constituents.

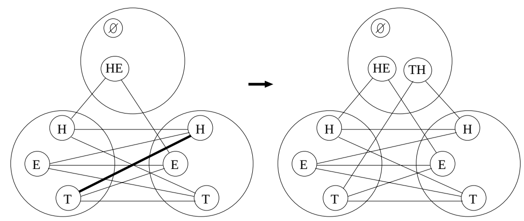

A novel use of Hebbian weight dynamics is the key to triggering chunking. The now 50-year-old Hebb rule says that the weight between two nodes should increase when the nodes are simultaneously active. Thus, the magnitude of the weight measures co-occurrences and can be used to directly signal chunking. This can be seen as automatically promoting 2nd-order correlations to become themselves first-class entities, allowing the subsequent detection of higher and higher order relationships, specifically higher order relationships that correspond to compositional structure.

The ability to grow the width of interactions captured by the model provides clear advantages over applying standard ML models in a fixed-width, sliding window fashion. Also, the ability to make predictions about any position provides more flexibility than predict-the-next-symbol paradigms. Space prevents discussion of the relations with HMMs and grammar induction methods.

In summary, this work can be seen as a long overdue continuation of work that makes predictions using hand-crafted CHs, such as early blackboard systems and IAM. We believe that compositional aggregation is one of the keys for bridging the gap between low-level, fine-grained representations and high level concepts, and that this is extremely important for long-lived autonomous systems in complex environments.

References

McClelland, J. L., and Rumelhart, D. E. 1981. An interactive activation model of context effects in letter perception: Part 1. an account of basic findings. *Psychological Review* 88:375–407.

Pfleger, K. 1998. Learning of compositional hierarchies for the modeling of context effects. Technical Report KSL–98–04, Stanford. See www-ksl.stanford.edu/publications/.