Procedia Computer Science
Volume 88, 2016, Pages 71–87
7th Annual International Conference on Biologically Inspired Cognitive Architectures, BICA 2016

Which Features Matter How Much When?

Lee Scheffler
Cognitivity, U.S.A.
www.cognitivity.technology
lee@cognitivity.technology

Abstract

How do brains learn which features matter how much, when and for what purposes? A specific feature may matter more or less for recognitions of different learned patterns, and in different contexts and attentional foci. Simple executable "neural circuits" built from biologically-inspired reusable memory pattern components in the NeurOS™ and NeuroBlocks™ technology[1] model and implement a range of learning and dynamic contextual/situational/attentional feature relevance. A pattern is a collection of weighted features, roughly analogous to a neuron or neuron assembly. New patterns are created for sufficiently novel feature combinations. Individual feature weights in best-matching existing patterns grow or diminish with repetition, yielding patterns that adjust to repeated experience. Arbitrarily complex classification meshes typical of human knowledge are easily assembled by varying a simple novelty parameter. Cascading pattern recognitions build up layers of concrete to abstract feature vocabularies. Names or labels are modeled as synonyms for experience patterns. Context can be modeled as yet another feature, derived from recent activity, to discriminate among otherwise similar patterns. Attention can be modeled as broad dynamic parameters modulating feature signal strengths.

Keywords: features, labels, learning, memory, cognition, context, attention, cognitive simulation

1 Introduction

An abiding challenge in understanding how biological brains work, and in creating similarly intelligent machines, is in understanding representations and algorithms of the information processing involved (Marr, 1982 pp. 19-27). How do biological brains achieve complex cognitive capabilities by interconnecting a relatively small variety of building blocks: neurons of several types, synapses and dendrites. How does useful learning start from a small number of data points and continually adjust with experience? How do different personal life experience histories yield similar-enough knowledge frameworks?

[1] patents pending; see www.cognitivity.technology

Selection and peer-review under responsibility of the Scientific Programme Committee of BICA 2016

doi:10.1016/j.procs.2016.07.409

This work follows Braitenberg's evolution-like "downhill invention" approach (Braitenberg, 1984 pp. 20-21): synthesize and incrementally improve working (artificial) systems that exhibit familiar cognitive capabilities. From these systems we may learn about core information representations and algorithms and system architectures underlying cognition. Insofar as these systems are built by interconnecting reusable biologically-inspired computational components in biologically plausible structures, we may draw some insights as to how corresponding capabilities may operate in biological brains, with rapid design iteration speeding hypothesis testing. And we may make strides along the path of creating artificial systems with similar capabilities. Put differently, let's start with simple assemblies of simple biologically-inspired components and see how far we can get.

What's a feature? For purposes of this paper, a feature is any distinct time-varying neural signal representing any percept or concept at any concrete through abstract level. This work explores how patterns of recurring feature combinations are learned, adjusted, classified, labeled, recognized and combined, and how recognition of such patterns vary with current context and attentional focus. Simple "neural circuits", using the NeurOS™ and NeuroBlocks™ technology, direct signals analogous to neuron spiking rates among reusable components performing biologically analogous computations and information storage. Small synthetic data sequence examples demonstrate the operation and learning dynamics of the assembled systems; these data are not "crafted" in any way other than to suggest one of many possible plausible sequences an organism might experience.

How can brains and intelligent systems learn so easily, quickly and effectively, especially from a few or even single examples, often with minimal or no supervision or reinforcement. How do such systems continually adjust to new experiences and environmental shifts? The bulk of current machine learning technologies require very large data sets, manual labeling, extensive parameterization of designs, take a long time to train and validate, and once trained are relatively static. Humans, by contrast, learn quickly from new examples and continually adjust what they know to new experience and corrections.

First, a single NeuroBlock "Set pattern" module managing an open collection of neuron-like feature patterns performs unsupervised learning of arbitrary classes. Learning creates new patterns or adjusts previous patterns. Varying matching and novelty parameters enables modeling of finely tuned exemplar patterns through increasingly broad stereotype patterns. Such a pattern is analogous to a "proxytype" (Prinz, 2002), where a learned pattern itself may be used as a feature input to other patterns. Feature inputs to a pattern can span multiple modalities and abstraction levels, depending on module connections.

Feeding the same input feature space in parallel to multiple such memory pattern modules with different matching and novelty parameter settings yields a rich classification mesh of learned exemplars and stereotypes that produces differential matching strengths to on-going experience.

Previously learned patterns are recognizable from a subset of their features, and their remembered features can merge with current input features to predict or "fill in the blanks" of missing features. Cascading pattern recognitions as features of other patterns yields layers of feature alphabets and vocabularies.

Names, labels and associated similar patterns are modeled as synonyms, a variation on a Set pattern with any/OR matching semantics: activity on any member feature activates the whole synonym pattern. Reification of such synonym patterns activates all synonyms, giving rapid access to similar patterns and labels.

Other cognitive processes (e.g., language recognition, other associational processes) may yield activation of a label, and, through reification of synonym patterns for the label, activate features associated with synonym patterns for the label, yielding a mechanism for imagination.

Interpretation of current inputs may depend on context. A word may have multiple meanings. The current context of a conversation, computed perhaps from recent words in a conversation, may disambiguate the word's meaning, elevating the activation of one candidate meaning over others.