Student Conference in Complexity Science

Book of Abstracts
Darwininan Neurodynamics

Brighton 19-22 August 2014

Darwininan Neurodynamics

Thursday, 21st August 2014

9:00 - 10:40 Darwinian Neurodynamics

Introduction to darwinian neurodynamics

Eörs Szathmáry, Eötvös Loránd University, Hungary

Online extreme evolutionary learning machines

Joshua Auerbach, École polytechnique fédérale de Lausanne, Switzerland (joshua.auerbach@epfl.ch)

Recently, the notion that the brain is fundamentally a prediction machine has gained traction within the cognitive science community. Consequently, the ability to learn accurate predictors from experience is crucial to creating intelligent robots. However, in order to make accurate predictions it is necessary to find appropriate data representations from which to learn. Finding such data representations or features is a fundamental challenge for machine learning. Often domain knowledge is employed to design useful features for specific problems, but learning representations in a domain independent manner is highly desirable. While many approaches for automatic feature extraction exist, they are often either computationally expensive or of marginal utility. On the other hand, methods such as Extreme Learning Machines (ELMs) have recently gained popularity as efficient and accurate model learners by employing large collections of fixed, random features. The computational efficiency of these approaches becomes particularly relevant when learning is done fully online, such as is the case for robots learning via their interactions with the world. Selectionist methods, which replace features offering low utility with random replacements, have been shown to produce efficient feature learning in one class of ELM. In recent work we have demonstrated that a Darwinian neurodynamic approach of feature replication can improve performance beyond selection alone, and may offer a path towards effective learning of predictive models in robotic agents.

Guiding Search with Neural Networks

Siddharth Sigtia, Queen Mary, University of London, UK (s.s.sigtia@qmul.ac.uk)

An algorithm is described that adaptively learns a non-linear mutation distribution. It works by training a denoising autoencoder (DA) online at each generation of a genetic algorithm to reconstruct a slowly decaying memory of the best genotypes so far. A compressed hidden layer forces the autoencoder to learn hidden features in the training set that can be used to accelerate search on novel problems with similar structure. Its output neurons define a probability distribution that we sample from to produce offspring solutions. The algorithm outperforms a canonical genetic algorithm on several combinatorial optimisation problems, e.g. multidimensional 0/1 knapsack problem, MAXSAT, HIFF, and on parameter optimisation problems, e.g. Rastrigin and Rosenbrock functions.

Information transfer is not enough to preserve systematicity

Emília Garcia Casdemont, Institut de Biologia Evolutiva Barcelona (IBE, UPF-CSIC), Spain (emilia.garcia.casademont@gmail.com)

Language involves a search process at two levels: (1) To process an utterance, a set of constructions needs to be found that is able to reconstruct the meaning of the utterance from the form (in parsing) or construct the form of the utterance from the meaning (in producing). Because there is often more than one construction that can trigger at any point in time (for example, the same word usually has multiple meanings) we get a search problem. (2) In addition to this, each language user has also the problem of finding which constructions are part of the shared language, based only on interactions with others. The language learner therefore unavoidably has to entertain several hypotheses, gathering enough evidence to decide which construction is part of the communal language. Here we focus on (2), more specifically on the role of information transfer within the framework of evolutionary neurodynamics.

We show in an agent-based experiment that the learning process, if seen as an evolutionary search process, requires information transfer between different hypotheses of communal constructions in order to explain why there is such systematicity in the language, for example, why the sequential pattern used in a noun-phrase with only an article and a noun is similar to one with an article, an adjective and a noun. We also show that it is not enough to achieve this information transfer but that also some form of ‚multi-level alignment' is needed to make the information transfer effective.

Symbol representation in the brain: past and current hypotheses

Yana Knight, Vrije Universiteit Brussel, Belgium (yana@ai.vub.ac.be)

The question of how the brain represents and manipulates symbol structures remains open. Studies have pointed to the existence of neural assemblies encoding a stimulus -- an idea that goes back to Hebb's associative learning and Abeles' synfire chains [1]. The neurons in the assembly are thought to be connected through temporal signals forming a spike train – neural code that may lie at the heart of encoding, decoding and information processing (id). Such assemblies may be permanent or even transient and exist as temporarily formed neural structures [2].

In 1988, Fodor and Pylyshyn [3] put forward a hypothesis that properties of human language and thought such as systematicity, productivity and compositionality suggest that the brain must implement a physical symbol system (PSS), formulated by Newell and Simon [4]. It has been pointed out that the non-symbolic mechanisms such as neural assembly and synfire chains may be used to implement and discover symbol systems. Fernando [5] presented a framework which illustrated how a PSS can be implemented in temporal coding. Within this framework, symbol tokens exist as spatiotemporal spikes on neuronal chains resembling synfire chains. This system is also capable of learning representations – an essential requirement for a realistic implementation of a symbol system.

Several attempts have focused on purely connectionist methods based on Hinton's distributed representations and allowed a network to form representations by statistically inferring constraints [6]. These, however, have succumbed to poor generalisation and training independence. Similarly, representing symbolic structures as semantic networks [7] has shown to suffer from the issue of rapid binding formation [5]. Forming part of a larger project dedicated to the investigation and implementation of linguistic structures in the brain, this paper surveys the existing symbol representation hypotheses and attempts to identify the future steps necessary to make progress with their implementation.

[1] Kumar, A., Rotter, S, Aertsen, A. 2010. Spiking activity propagation in neuronal networks: reconciling different perspectives on neural coding. Nature Reviews Neuroscience (11): 615-627. [2] Crick, F. 1984. Function of the thalamic reticular complex: The searchlight. In Proc. Natl. Acad. Sci, USA (81): 4586-4590. [3] Fodor, J. A., Pylyshyn, Z. W. 1988. Connectionism and cognitive architecture: a critical analysis. Cognition 28, 3-71. [4] Newell, A, Simon, H. A. 1976. Computer Science as Empirical Inquiry: Symbols and Search, „Communications of the ACM", Communications of the ACM 19 (3): 113–12. [5] Fernando, C. 2011. Symbol manipulation and rule learning in spiking neuronal networks, Journal of Theoretical Biology, Volume 275, Issue 1, 21 Pages 29-41. [6] Pollack, J.B. 1990. Recursive distributed representations. Artificial Intelligence 46(1), 77-105. [7] Marcus, G.F. 2001. The algebraic mind: integrating connectionism and cognitive science. MIT Press.