

Extending the Applicability of Neuroevolution

Giuseppe Cuccu

In recent years, machine learning has achieved results which were considered unthinkable just a decade ago. From object recognition in images to beating humans at *Go*, new results are introduced at a steady rate. Most come from a single, recent subfield of machine learning, namely deep learning, which trains deep (i.e. many-layer) neural networks using variations of the backpropagation algorithm. Backpropagation itself comes with a set of strict limitations and requirements however, which are arguably not warranted in a majority of applications. Alternative methods exist that sidestep such requirements entirely. One of these, considered the state-of-the-art training technique until the advent of deep learning, is *neuroevolution*. Neuroevolution trains neural networks using black-box optimization algorithms known as evolutionary algorithms, able to directly *search* in the space of neural networks. Extending the applicability of neuroevolution enables taking on classes of problems for which deep learning is unsuited.

This thesis studies and addresses some of the major issues currently limiting the applicability of neuroevolution, such as improving results on smaller networks, mitigating fitness stagnation, and scaling to large networks (i.e. higher dimensions). After an introduction to the key ideas and fundamental concepts it first makes the case for smaller, shallow networks, which can achieve high performance in complex tasks if the training data is carefully pre-processed. This has application in a wide variety of real-world problems which produce too little data and/or of too low quality to meet the requirements of deep learning methods. The first problem considered then is fitness stagnation, addressed through a novel restart strategy. Rather than terminating the run upon reaching convergence, the search is restarted on the most promising area of the space as derived from the history of the search so far. The thesis continues by taking on increasingly complex tasks with high dimensional observations. The feature extraction is separated from the decision making: with the former delegated to an external component, smaller networks are devoted entirely to decision making, highlighting their performance. Finally an evolutionary algorithm is presented to offer state-of-the-art performance while being specifically designed for neuroevolution. Its highly parallelizable implementation offers constant scaling over the size of the network, solely limited by the availability of parallel computational resources.

Jury:

Prof. Dr. Ulrich Ultes-Nietsche (president of the jury)

Prof. Dr. Philippe Cudr-Mauroux (thesis supervisor)

Prof. Dr. Denis Lalanne (internal examiner)

Prof. Dr. Julian Togelius (external examiner)