



Neuroevolution Algorithms

(Copyright © 2018 Piero Scaruffi - Silicon Valley Artificial Intelligence Research Institute)

By 2017 deep learning and reinforcement learning were the most popular techniques of A.I. but far from being accepted by everybody as the best. In fact, their very limitations led to a revival of evolutionary algorithms, especially when applied to evolving both the weights and the topologies of neural networks.

The first person to write about applying genetic algorithms to neural networks was the aerospace engineer Lawrence Fogel, a researcher at both UCLA and Convair, who published papers such as "Autonomous Automata" (1962) and "Toward Inductive Inference Automata" (1962), and then the book "Artificial Intelligence Through Simulated Evolution" (1966). The idea of training neural networks with genetic algorithm can also be found in John Holland's seminal book "Adaptation in Natural and Artificial Systems" (1975). The first attempts were carried out by Darrell Whitley at Colorado State University ("Applying Genetic Algorithms to Neural Net Learning", 1988); Lawrence Davis at BBN in Boston ("Mapping Classifier Systems into Neural Network", 1988); Rodney Brooks in person at MIT, who programmed a six-legged robot ("A Robot that Walks", 1989); Stanford students Geoffrey Miller (a future star of evolutionary psychology) and Peter Todd, who was also applying genetic algorithms to composing music ("Designing Neural Networks using Genetic Algorithms," 1989); Hugo de Garis at George Mason University ("Genetic Programming", 1990); Richard Belew, John McInerney and Nicol Schraudolph at UC San Diego ("Evolving Networks", 1990); David Schaffer, Rich Caruana and Larry Eshelman at Philips Laboratories in New York State ("Using Genetic Search to Exploit the Emergent Behavior of Neural Networks", 1990). That was the beginning of a new discipline, "neuroevolution". Initially, the research was limited to setting the weights of the network, i.e. to fixed-topology neuroevolution: generate a population of neural networks with their weights set randomly, measure the ones that are best at the task, let these ones generate a new population by mutating and crossing over with each other, and so on. Fixed-

topology algorithms were developed by Risto Miikkulainen's group at the University of Texas ("Evolving Finite State Behavior using Marker-based Genetic Encoding of Neural Networks", 1992); and by Randall Beer and John Gallagher at Case Western Reserve University ("Evolving Dynamical Neural Networks for Adaptive Behavior", 1992).

Later neuroevolutionists began to program their algorithms to also generate the topology of the network ("neurogenesis"): Eric Mjolsness at Los Alamos National Laboratory, who used simulated annealing (instead of genetic algorithms) to generate the structure of the network ("Scaling Machine Learning and Genetic Neural Nets", 1989); Steven Harp, Tariq Samad and Alope Guha of Honeywell labs in Minnesota, who developed NeuroGenesys ("Towards the Genetic Synthesis of Neural Networks", 1989); Jordan Pollack's GNARL (which stands for "GeNeralized Acquisition of Recurrent Links") at Ohio State University ("An Evolutionary Algorithm that Constructs Recurrent Neural Networks", 1994); Xin Yao's EPnet at Penn State ("A New Evolutionary System for Evolving Artificial Neural Networks", 1997); Josh Bongard's Artificial Ontogeny at the University of Zurich ("Evolving Complete Agents using Artificial Ontogeny", 2001); etc. In 2002 Risto Miikkulainen's student Ken Stanley at the University of Texas developed the algorithm NEAT (which stands for "NeuroEvolution of Augmenting Topologies"), the most popular and widely used algorithm for neuroevolution for a decade ("Evolving Neural Networks through Augmenting Topologies", 2002).

These were approaches of "direct encoding": they directly encoded network configurations. Then came the generation of "indirect encoding", that encodes a set of rules for generating networks, an approach pioneered by Hiroaki Kitano when he was at Carnegie Mellon University ("Designing Neural Networks using Genetic Algorithms with Graph Generation System", 1990), who was inspired by the work of the Hungarian biologist Aristid Lindenmayer when he was at City University of New York, who proposed a formal grammar called L-system to generate graphs ("Mathematical Models for Cellular Interactions in Development", 1968). Indirect encoding was popularized by the video titled "Evolved Virtual Creatures" (1994) of the virtual creatures "animated" by digital media artist Karl Sims' algorithms when he was artist-in-residence at Thinking Machines in Boston (in 2017 the video was at https://www.youtube.com/watch?v=JBgG_VSP7f8). A variant called "cellular encoding" was developed by Frederic Gruau in France ("Neural Network Synthesis using Cellular Encoding and the Genetic Algorithm", 1994).

The resurgence of evolution-based machine learning in the 2010s was centered around various Google laboratories: Quoc Le's team published "HyperNetworks" (2016), a variation on Stanley's HyperNEAT ("A Hypercube-Based Encoding for Evolving Large-Scale Neural Networks", 2009), while Daan Wierstra's team ("Convolution by Evolution", 2016) and Alex Kurakin's team ("Large-Scale Evolution of Image Classifiers", 2017) applied evolutionary methods to different kinds of deep learning.

In 2017 Ilya Sutskever's team at OpenAI announced a new algorithm that outperformed backpropagation, and it was simply an evolution (sorry for the pun) of the genetic algorithms of the 1970s ("Evolution Strategies as a Scalable Alternative to Reinforcement Learning").

A newly galvanized community of neuro-evolutionists seized the moment, and 2017 alone saw the introduction of DeepNEAT, a version of NEAT for deep networks, by Risto Miikkulainen, of NMODE (which stands for "Neuro-MODule Evolution") by Keyan Ghazi-Zahedi at the Max Planck Institute in Germany, and of "Genetic CNN" by Lingxi Xie and Alan Yuille at the Johns Hopkins University.

(Copyright © 2018 Piero Scaruffi - Silicon Valley Artificial Intelligence Research Institute)