



NOVA

IMS

Information
Management
School

MGI

Mestrado em Gestão de Informação

Master Program in Information Management

Impact of ALife Simulation of Darwinian and Lamarckian Evolutionary Theories

Yuliya Betkher

Dissertation presented as a partial requirement for the degree of Master of Information Management, specialization in Information Systems and Technologies Management

NOVA Information Management School
Instituto Superior de Estatística e Gestão de Informação
Universidade Nova de Lisboa

NOVA Information Management School
Instituto Superior de Estatística e Gestão de Informação
Universidade Nova de Lisboa

IMPACT OF ALIFE SIMULATION OF DARWINIAN AND LAMARCKIAN EVOLUTIONARY THEORIES

Yuliya Betkher

Dissertation presented as a partial requirement for the degree of Master of Information Management, specialization in Information Technologies and Systems Management

Advisor: Vitor Santos (NOVA Information Management School)

Co-advisor: Nuno Nabais (School of Arts and Humanities of the University of Lisbon)

September 2016

ABSTRACT

Until nowadays, the scientific community firmly rejected the Theory of Inheritance of Acquired Characteristics, a theory mostly associated with the name of Jean-Baptiste Lamarck (1774-1829). Though largely dismissed when applied to biological organisms, this theory found its place in a young discipline called Artificial Life. Based on the two abstract models of Darwinian and Lamarckian evolutionary theories built using neural networks and genetic algorithms, this research aims to present a notion of the potential impact of implementation of Lamarckian knowledge inheritance across disciplines. In order to obtain our results, we conducted a focus group discussion between experts in biology, computer science and philosophy, and used their opinions as qualitative data in our research. As a result of completing the above procedure, we have found some implications of such implementation in each mentioned discipline. In synthetic biology, this means that we would engineer organisms precisely up to our specific needs. At the moment, we can think of better drugs, greener fuels and dramatic changes in chemical industry. In computer science, Lamarckian evolutionary algorithms have been used for quite some years, and quite successfully. However, their application in strong ALife can only be approximated based on the existing roadmaps of futurists. In philosophy, creating artificial life seems consistent with nature and even God, if there is one. At the same time, this implementation may contradict the concept of free will, which is defined as the capacity for an agent to make choices in which the outcome has not been determined by past events. This study has certain limitations, which means that larger focus group and more prepared participants would provide more precise results.

KEYWORDS

Artificial Life; Philosophy; Genetic Algorithms; Darwinism; Lamarckism.

PUBLICATIONS RESULTING FROM THIS DISSERTATION

Betkher, Y., Nabais, N., & Santos, V. (2016). Impact of ALife Simulation of Darwinian and Lamarckian Evolutionary Theories (L. George, H. K. Singh, & S. Elsayed, Eds.). The 20th Asia-Pacific Symposium on Intelligent and Evolutionary Systems (IES 2016).

Betkher, Y., Nabais, N., & Santos, V. (2016). Impact of ALife Simulation of Darwin's and Lamarck's Theories of Evolution on Life and Society (C. Gershenson, T. Froese, J. M. Siqueiros, W. Aguilar, E. Izquierdo, & H. Sayama, Eds.) [Abstract]. Late Breaking Abstracts Booklet Artificial Life Conference 2016, 13-14.

TABLE OF CONTENTS

1. Introduction	1
1.1. Background and problem identification	2
1.2. Study relevance and justification	3
1.2.1. Current application in engineering	4
1.2.2. Some later findings in computational biology	5
1.2.3. New significance in philosophy and ethics	5
1.3. Research objectives	6
2. Theoretical framework	7
2.1. Evolution	7
2.2. Genetic algorithms	7
2.3. Artificial neural networks and backpropagation algorithm	8
2.4. Artificial life and its philosophical meaning	12
3. Methodology (Design science)	15
3.1. Design science	15
3.2. Focus group	15
3.3. Research strategy	16
4. Research	17
4.1. Fundamentals	17
4.2. Proposal (model)	17
4.3. Preliminary conclusions	22
4.4. Validation	23
5. Discussion	25
5.1. Biological impact	25
5.2. Technical impact	26
5.3. Philosophical impact	27
5.4. Analyzing focus group data	28
6. Conclusions and future work	30
6.1. Main contributions	30
6.2. Limitations of the current work	31
6.3. Future work	31
7. Bibliography	33

1. INTRODUCTION

The behavior of natural organisms in the real world is not fixed across their lifespan. Through interactions with the environment, they gain experience and develop a tendency to repeat the actions that bring pleasure or benefit, and to avoid those that lead to danger or pain.

At the same time, organisms are not born in a blank state – they develop according to the information in their genes, which are inherited from the ancestors and selected through the struggle for existence (Sasaki and Tokoro, 2000).

However, only what is inborn in their own heritage can be transferred along with their genes. The acquired characters will not be encoded in the genes, and therefore will not be directly passed to the offspring – according to *Darwinism*, all the knowledge that biological organisms have gained should be developed over again by each new generation.

Here a simple question arises: what if living beings could get past this limitation? What if children could pick up where parents left off developing their expertise, health, coordination and reflexes, each generation building on the last to reach out for higher and higher goals?

This attractive, although hypothetic process is called *Lamarckian inheritance*, a long-discredited mechanism of evolution. Through learning, individuals would experience certain adaptive changes and acquire new traits that would be directly transmitted to their offspring. Although it was largely dismissed as a valid theory for natural systems, Lamarckian evolution found its place and proven effective within computer applications (Ross, 1999).

But science moves ahead, and in the nearest future, the implementation of Lamarckian evolution may turn into something bigger than just a cybernetic adventure. Even though our commonly accepted definition of life does not yet recognize any current simulations or applications as alive, it may not always stay that way. The opinions regarding this matter vary, but according to the *strong ALife* position, first introduced by Neumann in 1963, life can be abstracted away from any particular medium.

In this research, we are going to evaluate the potential impact of the implementation of Lamarckian evolution, in particular, of the inheritance of skills and knowledge, given the possibility of creating life within computational environment. Instead of being just an engineering problem, it becomes a cross-disciplinary topic that creates numerous philosophical questions and implications.

In order to demonstrate that it is possible to implement Lamarckian evolution in a computational environment, we refer to previously created models and present our own, where a *neural network* is regarded as a learnable individual (Rumelhart and McClelland, 1986), and *genetic algorithms* (Holland, 1975) are applied to the population of such individuals based on mechanisms of natural evolutionary processes and genetics. Using a focus group of researchers from computer science, biology and philosophy, we have validated the model and evaluated its potential impact on different matters in our lives, including technology, ethics, life and society.

1.1. BACKGROUND AND PROBLEM IDENTIFICATION

Jean-Baptiste Lamarck (1744-1829) and Charles Darwin (1809-1882) both contemplated and developed ideas about how life on earth evolved to be the way it is now. They both believed that living things change to be better fit and adapted to their environments, and that all organisms relate to one another.

Lamarck's *Theory of Inheritance of Acquired Characteristics*, first described in his work "Philosophie Zoologique" in 1809, implied that the adaptive changes that species may undergo through interactions with the environment are passed on to their offspring and later generations. In the classic example, the giraffe obtained its long neck by stretching to reach higher branches. This stretching experience was further transmitted to future generations with each getting a slightly longer neck. In other words, the species was being directly changed by its interaction with the environment.

Darwin, on the contrary, in his *Theory of Evolution*, published in his book "On the Origin of Species" in 1859, stated that the offspring are born with their parents' beneficial traits, and as they reproduce, individuals with that trait make up more of the population, while less adapted individuals die off. He introduced a plausible mechanism called *natural selection*, that acts to preserve and accumulate minor advantageous genetic mutations. If a specimen develops a functional advantage, e.g. grows wings and learns to fly, its offspring would inherit that advantage and pass it on to their offspring. The inferior (disadvantaged) members of the same species would gradually die out, leaving only the superior (advantaged) members of the species (Sharma, 2010).

Darwin's Theory of Evolution has been supported by evidence from a wide variety of scientific disciplines, and Lamarck's Theory of Inheritance of Acquired Characteristics has been proven wrong, even though modern research in the emerging field of genetics called *epigenetics* has shown that Lamarck may have been at least partially correct all along (Springer and Holley, 2013). The mainstream of modern evolutionary theory follows Darwinism and denies the possibility of direct inheritance of acquired traits (Sasaki and Tokoro, 2000).

Nevertheless, artificial organisms would reproduce in a completely different manner than biological, under at least partial conscious control, giving it a Lamarckian component (Farmer and Belin, 1990). Unlike life in the natural world, computer programs use uncomplicated transformations between genotypes and phenotypes, and the inversion of phenotypes to their corresponding genotypes is often manageable. In certain cases, where genotypes are their own phenotypes, no transformation is needed at all. The significance of this with respect to Lamarckian evolution is that it is possible to optimize a phenotype in the context of a particular problem environment, and have this optimization represented in a corresponding genotype for succeeding inheritance by offspring (Ross, 1999).

Considering the technology trends and the history of humanity, our evolution provides evidence that humans will one day create machines more intelligent than they are. The reasons to believe in the creation of a conscious machine include the exponential increase of computational capacity of computers, automatic knowledge acquisition and algorithms like recursion, neural networks, and genetic algorithms (Kurzweil, 1999). Kurzweil predicts the machines will appear to have their own free will and even spiritual experiences.

If the inception of strong AI and strong ALife is near, the implementation of Lamarckian inheritance may impact our lives in different domains, from technological applications to philosophical and ethical concerns. The problem being addressed here is how exactly this implementation can influence various disciplines and what it can potentially lead to. In this research, we will give a special attention to the possible future role of knowledge inheritance by the Lamarckian scheme in engineering, biology, philosophy and ethics.

1.2. STUDY RELEVANCE AND JUSTIFICATION

Throughout the history, scientists have studied evolution for the same reasons that they have learned any other discipline — the thirst for knowledge, the desire to understand the past and predict the future, and the necessity to organize our world. Evolution, especially the understanding of how organisms evolve through natural selection, has always been an area of science with various practical applications (Bull and Wichman, 2001).

But nowadays, with the emergence of artificial intelligence, computational neuroscience and transhumanism, evolutionary studies have acquired fundamentally different ethical and social significance that extends beyond simple curiosity. Already existing roadmap on *whole brain emulation* (Sandberg and Bostrom, 2008) attempts to achieve software intelligence by copying the function of biological nervous systems into software. This approach produces numerous ethical issues that should affect responsible policy for developing the field. Animal emulations have controversial moral status, and a principle of analogy is suggested for judging treatment of virtual animals. Various considerations of developing and utilizing human brain emulations are discussed (Sandberg, 2014).

Among the latest published books, taking inspiration from self-awareness in humans, the new notion of computational self-awareness as a fundamental concept for designing and operating computing systems has been introduced (Lewis et al., 2016). The basic ability of such self-aware computing systems is to gather information about their state and progress, learning and maintaining models containing knowledge that enables them to reason about their behavior. Self-aware computing systems will have the ability to utilize this knowledge to effectively and autonomously adapt and explain their behavior in dynamic environments.

Although the accuracy of predictions of future developments in AI and ALife is difficult to evaluate, according to Ray Kurzweil himself, 89 out of 108 predictions he made so far were entirely correct by the end of 2009. An additional 13 were what he calls “essentially correct” (meaning that they were likely to be realized within a few years of 2009), for a total of 102 out of 108. Another 3 are partially correct, 2 look like they are about 10 years off, and 1, which was tongue in cheek anyway, was just wrong (Wang, 2010).

While at present, whole brain emulation seems an unfeasibly ambitious challenge, the necessary computing power and various scanning methods are rapidly developing. Large-scale computational brain models are a very active research area, at present reaching the size of mammalian nervous systems (Djurfeldt et al., 2008; Eliasmith et al., 2012; Markram, 2006; Preissl et

al., 2012). Whole brain emulation can be considered the logical endpoint of current trends in computational neuroscience and systems biology (Sandberg, 2014).

The implementation of Lamarckian evolution in future ALife systems such as virtual lab animals has wide practical application across disciplines. We have examined the trends and role of such evolution in engineering, computational biology and ethics, and based on this analysis, came to several assumptions that are to be evaluated by the focus group.

1.2.1. Current application in engineering

In a dynamic and unpredictable environment such as real world, it is very difficult to construct intelligent machines or computer programs that would perfectly manage to produce desirable results from the very beginning. Therefore, an approach based on adaptive computation or evolutionary computation, where programs adapt themselves towards given situations through generating and testing, gained its popularity and significance (Sasaki and Tokoro, 2000).

Simulations of evolution using evolutionary algorithms originate from the work of Barricelli in the 1960s, continued by Fraser, who published a series of papers on simulation of artificial selection (Fraser, 1958). As a result of the work of Rechenberg, who used evolution strategies in the 1960s and early 1970s to solve complex engineering problems, artificial evolution became a widely recognized optimization method (Rechenberg, 1973). Genetic algorithms in particular became well-known through the writing of Holland (1975). As academic interest grew, dramatic increases in the power of computers allowed practical applications, including the automatic evolution of software (Koza, 1992). Evolutionary algorithms are now applied in solving multi-dimensional problems more efficiently than computer programs developed by human designers, and also to optimize the design of systems (Jamshidi, 2003).

A research area called *artificial life* (Langton, 1989) is a typical example that analyzes mathematical aspects of the dynamics residing in life in a synthetic way and tries to apply principles of natural systems (ranging from swarms of cells to human societies) as models for possible novel methods of adaptive computation. In software-based artificial life, neural networks are often applied in modeling the brain of an agent. Although traditionally more of an artificial intelligence technique, neural nets can be used for simulating population dynamics of organisms with an ability to learn (Kumar and Bhatnagar, 2010). Genetic algorithms are applied to such populations based on evolutionary and genetic mechanisms.

Due to the biological background, earlier attempts of artificial life modeling have always focused on a Darwinian evolution, based on competition of artificial beings in a computational environment, where new artificial organisms would appear only as the result of combining morphology of parents (Hinton and Nowlan, 1987). Until nowadays, Darwin's evolutionary models have been widely used in different scientific fields. Many of such implementations were motivated by the idea of constructing practical devices that have some of the useful features of living systems, such as robustness, flexibility, and autonomy (Bedau, 2002).

At the same time, from the engineering point of view, it is not necessary to consider only Darwinian models. The possibility of heredity of acquired characteristics can be quite useful, and

several studies have already shown the significant increase in performance of problem-solving systems using Lamarckian scheme (Grefenstette et al., 1990; Davidor, 1991).

In evolutionary algorithms, the implementation of Lamarckian inheritance means that an individual can modify its genetic code during or after fitness evaluation, or lifetime. This idea has been used in several studies with particular success in problems where the application of a local search operator obtains a substantial improvement, e.g. traveling salesman problem (Ross, 1999). The effectiveness and superiority of Lamarckian evolutionary algorithm has also been demonstrated for fixed tasks in stationary environments, even though Darwinian population adapts better to dynamic environments (Sasaki and Tokoro, 2000).

1.2.2. Some later findings in computational biology

Relatively recent research in cell biology has shown that the internal chemistry of living cells is a form of computation (Bray, 2011). Such ideas are currently breaking boundaries between scientific disciplines and give rise to interdisciplinary sciences like *computational biology*, which involves the development and application of data-analytical and theoretical methods, mathematical modeling and computational simulation techniques to the study of biological, behavioral, and social systems (Huerta et al., 2000).

A *wetware computer* is an organic computer (also known as an artificial organic brain or a neurocomputer) built from living neurons. Professor Ditto, at the Georgia Institute of Technology, is the primary researcher driving the creation of these artificially constructed, but still organic brains. One prototype is constructed from *leech neurons*, and is capable of performing simple arithmetic operations. The concepts are still being researched and prototyped, but in the near future, it is expected that artificially constructed organic brains, even though they are still considerably simpler in design than animal brains, should be capable of simple pattern recognition tasks such as handwriting recognition (Borresen and Lynch, 2009).

At the same time, while originally dismissed as non-feasible, Lamarckian evolution now appears more and more in biological systems ranging from microbes to mammals, and molecular mechanisms that might realize this mode of inheritance are being clarified. Epigenetics, a set of means to propagate a phenotypic change across generations, appears to provide a set of feasible molecular means that may realize Lamarckism. In addition, several mechanisms exist which may allow the phenotype to instruct the genotype at a given environment. Recent advances in molecular evolution have been surveyed and realistic means have been presented to engineer Lamarckian organisms in the lab which might possess improved evolvability (Pilpel, 2016).

1.2.3. New significance in philosophy and ethics

Computational biology gives one the sense that we are at the threshold of yet another of civilization's "Spinoza moments" where the entire framework for thinking about life is dramatically, and irrevocably restructured. The idea that cellular membranes and contents may be functional equivalents of computers does not appear strange and implausible any longer. And even if the

implementation of strong ALife is a matter of future, considering potential risks and their ethical impacts is an important aspect of research ethics, even when dealing with merely possible future radical technologies (Sandberg, 2014).

Evolutionary studies have provided us better understanding of ourselves and helped us find our own place on Earth with 1.8 million identified species, and possibly 10 million total species. The context of evolution gives an insight on how to behave among members of our own and other species. Evolution helps us understand the purpose and reasons for our physiology and anatomy (Moritz, 2010).

Since the Darwinian theory of evolution gained widespread acceptance in the late 1800s, scientists and philosophers have been looking for ways to relate traditional evolutionary theory to the way we live, interact with society, and think about our place in existence. Now the Lamarckian evolution within artificial life has become a relatively recent object worthy of philosophical attention (Stewart, 2005). Therefore, the new questions of particular interest in evolutionary philosophy are how much of an influence Lamarckian evolution in ALife would have on human behavior, and what are the philosophical implications of this evolution on issues that relate to ethics and morality.

1.3. RESEARCH OBJECTIVES

The main goal of this dissertation is to identify the potential impact of implementation of knowledge inheritance in artificial organisms using Lamarckian scheme across disciplines. The objectives being pursued in order to achieve this goal are the following:

1. Build two artificial life models of Darwinian and Lamarckian knowledge inheritance processes using genetic algorithms and artificial neural networks. Based on this example, prove that such implementation is possible in computational environment.
2. Evaluate possible impact of knowledge inheritance in artificial organisms on life and society, considering latest trends across disciplines.
3. Using the qualitative data obtained from focus group discussion, evaluate and understand the impact of Lamarckian evolution in ALife in computer science, biology and philosophy.

2. THEORETICAL FRAMEWORK

2.1. EVOLUTION

Evolution is variation in the heritable traits of biological populations over succeeding generations. Evolution is the mechanism producing the diversity of life at every level of biological organization, including the levels of species, individual organisms, and molecules (Hall et al., 2007).

Evolution is the foundation of modern science, believed as one of the most reliably proven by all facts and theories of science, based on evidence not just from the biological sciences, but also from anthropology, psychology, astrophysics, chemistry, geology, physics, mathematics, and other scientific disciplines, as well as behavioral and social sciences. The discovery of evolution has made significant contributions to humanity, including the prevention and treatment of human disease, new agricultural products, industrial innovations, a subfield of computer science, and rapid advances in life science (Ayala, 2008).

Evolution and philosophy relate to each other as long as the idea of evolution exists. On the one hand, this is due to the fact that science and philosophy only separated around the time evolutionary theories were being first proposed, on the other hand, because evolution was opposed to many cherished philosophical doctrines, particularly in Darwinian context (Wilkins, 1997).

2.2. GENETIC ALGORITHMS

A *genetic algorithm* (or GA for short) is a programming technique that imitates biological evolution as a problem-solving strategy. Given a specific problem to solve, the input to the GA is a set of candidate solutions to that problem, encoded in a certain way, and a metric called a fitness function that quantitatively evaluates each possible solution. These candidates can possibly be already known working solutions, with the purpose of the GA being to improve them, but more often they are randomly generated (Marczyk, 2004).

Genetic algorithms (GAs) provide a learning method inspired by an analogy to biological evolution. Rather than search from general-to-specific hypotheses, or from simple-to-complex, GAs generate successor hypotheses by repeatedly mutating and recombining parts of the best currently known hypotheses. At each step, a set of hypotheses named the current population is renewed by replacing some part of the population by the offspring of the fittest to current hypotheses. The process constitutes a generate-and-test beam-search of hypotheses, in which variations of the best current hypotheses are most likely to be considered next (Mitchell, 1997).

In GAs, the term chromosome normally refers to a point in the search space of candidate solutions to a problem, often encoded as a bit string. The genes are either single bits or short blocks of adjacent bits that encode a certain element of the candidate solution. Each locus in the chromosome has two possible alleles: 0 and 1; for larger alphabets other alleles are possible (Mitchell, 1996).

The most generic form of a genetic algorithm involves three types of operators: selection, crossover (single point), and mutation. Selection picks chromosomes in the population for

reproduction: the fitter the chromosome, the more likely it will be selected. Crossover is typically the commutation of genetic material between two single chromosome haploid parents. Mutation flips the bit at a randomly chosen locus (or, for larger alphabets, replaces the symbol at a randomly chosen locus with a randomly chosen new symbol); it occurs with a very small probability (e.g. 0.001).

The GA processes populations of chromosomes, consequentially replacing one such population with another. The fitness function assigns a score (fitness) to each chromosome in the current population. The fitness of a chromosome depends on how well that chromosome solves the given problem.

The *genotype* of an individual in a GA using bit strings is simply the configuration of bits in that individual's chromosome. Even though there is often no such thing as a phenotype in the context of GAs, lately some researchers have experimented with GAs in which there is both a genotypic level and a phenotypic level, e.g. the bit-string encoding of a neural network and the neural network itself.

Genetic algorithms are widely applied in different scientific areas like electronics, mechanics and computer science. These are just a tiny sample of their possible applications, although those are some of the most important, and many other examples can be found in advanced books and articles.

2.3. ARTIFICIAL NEURAL NETWORKS AND BACKPROPAGATION ALGORITHM

One of the ways to understand and resolve complex problems is to follow the lemma “divide and conquer” and decompose them into simpler elements. Also simple elements may be assembled to produce a complex system (Bar Yam, 1997). Networks are one method for achieving this. There is a large number of various types of networks, but they all are characterized by the following components: a set of nodes, and connections between nodes (Gershenson, 2003).

Stergiou and Siganos (1996) defined *artificial neural network* (ANN) as an information processing paradigm inspired by the way biological nervous systems (such as the brain) process information. It consists of interconnected processing elements working together to solve specific problems and learns by example: each ANN is built for a specific application through a “learning process”. In fact, ANNs have also inherited another important characteristic of a brain: the ability to interpolate from incomplete information (Hewitson and Crane, 1994).

An *artificial neuron* is a computational model inspired by biological neurons. The interface through which biological neurons interact with their neighbors usually consists of several axon terminals connected via synapses to dendrites on other neurons. If the sum of the input signals into one neuron exceeds a certain threshold, the neuron sends an action potential (AP) at the axon hillock and passes this electrical signal along the axon (Weiss, 2007). This signal might be sent to another synapse, and might activate other neurons (Gershenson, 2003).

The complexity of real neurons is drastically simplified when modeling artificial neurons, which consist of *inputs* (like synapses) multiplied by *weights* (strength of the respective signals), and then are computed by a mathematical function which determines the activation of the neuron.

Another function (which may be the identity) computes the output of the artificial neuron (sometimes in dependence of a certain threshold). ANNs connect artificial neurons in order to process information.

A neural network can be characterized by its *architecture* (which is pattern of connections between neurons), by *algorithm* (or method of determining weights on the connections), and by its *activation function* (Fausett, 1994).

The arrangement of neurons into layers and the connection patterns within and between layers is called the *net architecture*. Neural nets are often classified as *single-layer* or *multilayer* (Fausett, 1994). When defining the number of layers, the input units are not counted as a layer, since they do not perform any computation. Therefore, the number of layers in the net can be determined by the number of weighted interconnected links between the slabs of neurons, as those weights in the network contain extremely important information.

A single-layer neural network has one layer of connection weights (see figure 1). The input units receive signals from outside world, while output units represent the response of the net. In this typical single-layer neural network, the input units are fully connected to output units, while being unconnected to one another; the output units are not connected to each other either.




Figure 1: A single-layer neural net

The training in this neural net setting is typically accomplished by presenting a sequence of training vectors, each with an associated output vector. The weights are being changed according to a learning algorithm, which is called *supervised training*. The single-layer nets use supervised training (the *Hebb rule* or the *delta rule*). *Unsupervised* or *self-organizing* nets group similar input vectors together without the use of training data to identify to which group each vector belongs, and are often used for clustering problems, as well as for other tasks. There is some ambiguity in the labeling of training methods as supervised or unsupervised, and some authors find a third category, called *self-supervised training*, useful.

The characteristics of the problem to be resolved determines whether a single-layer net is adequate. In case the problem is more difficult, a multilayer net, such as that trained by *backpropagation*, may be better. A multilayer neural net (see figure 2) is a net with one or more

layers (or levels) of nodes (the so-called hidden units) between the input units and output units. Normally, there is a layer of weights between two neighbor levels of units (input, hidden, or output). Although training multilayer neural nets may be more difficult, they can resolve problems that single-layer nets cannot be trained to perform correctly at all.




Figure 2: A multilayer neural net

The single-layer and multilayer networks illustrated in figures 1 and 2 are examples of *feedforward* nets, where the signals flow from the input units to the output units, in a forward direction. The fully interconnected *competitive* net in the figure 3 is an example of a *recurrent* net. In this network, there are closed-loop signal paths from a unit back to itself.




Figure 3: A competitive neural net

The *backpropagation algorithm* (Rumelhart and McClelland, 1986) or the *generalized delta rule* (Fausett, 1994) is applied in layered feed-forward ANNs, where artificial neurons are ordered in

layers and send their signals “forward”, and then the errors are propagated backwards. It is simply a gradient descent method of minimizing the total squared error of the output computed by the net. The network receives inputs by neurons in the input layer, and the output of the network is given by the neurons on an output layer. There may be one or more intermediate hidden layers.

The backpropagation algorithm uses supervised learning, which means that we initially supply the algorithm with examples of the inputs and outputs we want the network to compute, and then calculate the error (difference between actual and expected results). The goal of the backpropagation algorithm is to reduce this error, until the ANN learns the training data (Gershenson, 2003). The training begins with random weights, and the objective is to adjust them so that the error will be minimal.

The training of the network by backpropagation involves three stages: the feedforward of the input training pattern (i), the calculation and backpropagation of the associated error (ii), and the adjustment of the weights (iii). While a single-layer net is very limited in the mappings it can learn, a multilayer net can learn any continuous mapping to an arbitrary accuracy. Usually one hidden layer is sufficient, even though more layers may be beneficial for some applications (Fausett, 1994).

An activation function for a backpropagation net should have several important characteristics: *continuity*, *differentiability* (a derivative should exist at each point in its domain), it should be *monotonically non-decreasing* and easy to compute. One of the most typical activation functions is the binary sigmoid function, which has range of (0, 1) and is defined as:

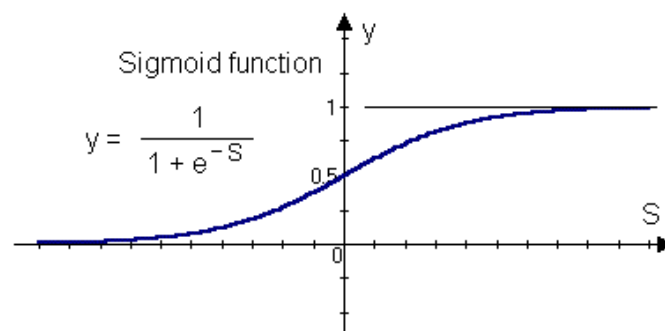


Figure 4: A binary sigmoid function

The choice of initial weights and biases will influence whether the net reaches a global (or only a local) minimum of the error and, if so, how quickly it converges. The update of the weight between two units depends on both derivative of the upper unit’s activation function and the activation of the lower unit. Therefore, it is crucial to avoid choices of initial weights that would make it likely that either activations or derivatives of activations are zero.

A typical procedure is to initialize the weights (and biases) to random values between -0.5 and 0.5 (or between -1 and 1 or some other required interval). The values can be positive or negative, since the final weights after training can be of either sign as well. It is important to remember that the values for the initial weights should not be very large, otherwise the initial input signals to each hidden or output unit will be likely to fall in the region where the derivative of the sigmoid function has a very small value (the so-called *saturation region*). At the same time, if the

initial weights are too small, the net input to a hidden or output unit will be almost zero, which also causes extremely slow learning.

A simple modification of random initialization, developed by Nguyen and Widrow (1990) typically gives much faster learning. Their approach is based on a geometrical analysis of the response of the hidden neurons to a single input. The Nguyen-Widrow method generates initial weights and bias values for a layer, so that the active regions of the layer's neurons will be distributed approximately evenly over the input space (Demuth, 2002).

A *Darwin-neural network* is a neural network based on structural patterns and learning processes introduced by the neural Darwinism theory (Edelman, 1987). A Darwin-neural network learns specific tasks through interactions with an unknown environment, and its behavior develops according to gained the experience (Manderick, 1991). However, a *Lamarckian-neural network* is a non-orthodox problem solving tool that combines evolution and learning techniques (Cortez et al., 2002).

2.4. ARTIFICIAL LIFE AND ITS PHILOSOPHICAL MEANING

The contemporary idea of *artificial life* as a discipline was proposed by Christopher Langton (1989) in his book "Artificial Life. An Overview", who described the field broadly as devoted to studying the scientific, technological, artistic, philosophical, and social implications of creating "living" artifacts. According to Langton, it emulates traditional biology by trying to recreate some features of biological phenomena; at the same time, the modeling philosophy of artificial life strongly differs from traditional modeling by studying not only "life-as-we-know-it" but also "life-as-it-might-be". In addition, as Bruce Sterling later wrote for "The Magazine of Fantasy and Science Fiction" in 1992, the relatively new field of study named artificial life was created as an attempt to abstract the logical form of life from its material manifestation.

The concept of artificial life may be used with different meanings. At the beginning, the term "artificial life" was originally defined it as "life made by man rather than by nature," i.e., it is the study of man-made systems that exhibit behaviors characteristic of natural living systems (Langton, 1989). Later, however, Langton discovered fundamental problems with this definition, and changed it to "the study of natural life, where nature is understood to include rather than to exclude, human beings and their artifacts" (Langton, 1998). He insisted that human beings and all their actions are inseparable part of nature, therefore a major goal of ALife should be to work towards eliminating "artificial life" as a phrase that differs in meaning in any significant way from the term "biology". In fact, it is quite common nowadays in a biological research to use computational models, which would have been considered ALife 20 years ago, but now they are part of mainstream biology (Bourne et al., 2005).

Since its outbreak in the late 1980s, the study of artificial life has bonded together scientists interested in a formal and general comprehension of living systems (Husbands et al., 1997). Researchers and philosophers are actively conducting their studies in the areas like the definition of life, the relationship between the life and mind, and the possibility of creating life within computational environment (Keeley, 1998).

Many of the subjects of artificial life and artificial intelligence overlap, so these fields can be considered closely related. According to Bedau (2003), living and evolving in a dynamic and unpredictable environment requires at least rudimentary intelligence. Nevertheless, artificial life is especially focused on systems, which can mimic nature and its laws and therefore it is more related to biology, while the latter is mainly focused on how human intelligence can be replicated, and therefore, it is more related to psychology. In addition, different modeling strategies are used in these two fields. Most conventional AI models are top-down specific systems involving a complicated, centralized controller that makes decisions based on access to all aspects of global state. However, ALife systems are usually bottom-up (Maes, 1993), implemented as low-level agents that simultaneously interact with each other, and whose decisions are based on information about, and directly affect, only their own local environment (Bedau, 2003).

The inception of artificial life has its deep philosophical meaning. It makes humans rethink their conventional anthropocentric views and triggers multiple questions about nature and meaning of life. It aims at understanding the fundamental behavior of life-like systems by synthesizing that behavior in artificial systems.

Philosophy and artificial life are true intellectual partners for many reasons. According to Sellars (1963), the purpose of philosophy, abstractly formulated, is to understand how things in the broadest possible sense of the term hang together in the broadest possible sense of the term. Bedau concluded in 2002, that it applies to artificial life as well: both seek to acquire understanding of phenomena at a level of generality that is sufficient to ignore contingencies and uncover essential natures.

Even though philosophy cannot claim with certainty what is the real answer to the doubts it raises, it is able to offer many opportunities that expand our thoughts; therefore, while reducing the feeling of certainty as to what things are, it significantly improves the knowledge as to what they may be (Russell, 1959). Artificial life simulations attempt to answer similar "What if X?" questions (Bedau, 2002), but the premises they represent are complicated enough to be researched only by computer simulation.

Artificial life is an interdisciplinary field, not just a scientific and engineering enterprise, as Bedau noted in "The scientific and philosophical scope of artificial life" (to appear in "Leonardo") in 2002. He added that since artificial life offers a new prospect on the primary nature of many basic aspects of reality like life, adaptation, and creation, it has rich implications for a handful of broad philosophical topics: emergence, evolution, life, and mind.

At present, the commonly accepted understanding of life does not acknowledge any current ALife simulations or software to be alive, as well as they do not form part of the evolutionary process of any ecosystem. However, different views on artificial life's potential have arisen (cf. Strong AI vs. Weak AI).

Weak ALife position does not accept the possibility of generating a "living process" outside of a biological solution. The researchers in weak ALife try to simulate life processes to understand the underlying mechanisms of biological phenomena.

Strong ALife amounts to the claim that, by programming a computer, one can literally bring bits of its hardware to life (Olson, 1997). John von Neumann described strong ALife position through his Theory of Reproducing Automata (1963), which states that life is a process which can be abstracted away from any particular medium. Several decades later, Ray T. declared that his program Tierra is not simulating life in a computational environment, but generating it, as it was noted in C. Taylor's "To follow a rule" in 1992.

In this research, the role of ALife system modeling is to serve to prime intuitions and generate hypotheses regarding living systems, which may later be tested by more traditional means. On this view, we are talking about weak ALife, and the modeling and simulation of evolution may best be thought of as a theory building paradigm (Diallo et al., 2013).

3. METHODOLOGY (DESIGN SCIENCE)

3.1. DESIGN SCIENCE

Most of the research in the Information Systems discipline can be characterized by two paradigms: *behavioral science* and *design science* (Hevner et al., 2004). Considering the nature and the objectives of this dissertation, design science was chosen as a suitable methodology: it is outcome-based, and has explicit intention of improving the functional performance of the model (Vaishnavi et al., 2007).

The design-science paradigm originates from engineering and the sciences of the artificial (Simon, 1996). It helps to understand the behavior of information systems by creating new and innovative artifacts (Hevner et al., 2004). Such artifacts, depending on the research, are widely defined by constructs, models, methods and instantiations. In design science, as opposed to explanatory science, academic research can be seen as a quest for understanding and improving human performance (Van Aken, 2005), which constitutes one of the possible applications of artificial life as a discipline and this research in particular.

Hevner et al. (2004) have introduced a set of guidelines for design science research within the discipline of Information Systems. In this dissertation (following the guidelines), an abstract evolutionary model will be considered an artifact. The artifact will be validated by a focus group.

3.2. FOCUS GROUP

A *focus group* is a form of qualitative research where a group of people are asked about their perceptions, opinions, beliefs and attitudes towards a certain product, service, concept, advertisement, idea, or packaging. A moderator asks questions in an interactive group setting where participants can talk and engage in a discussion with other group members. During this process, the researcher either takes notes or makes a record of the crucial points he or she is getting from the group. Care should be taken while selecting members of the group to obtain effective and authoritative responses (Morgan, 1997).

Group discussion produces data and insights that would be hardly accessible without interaction found in a group setting—hearing others' outspoken experiences triggers associations, ideas and memories in participants. This is also known as the group effect where group members engage in some sort of 'chaining' or 'cascading' effect – new ideas emerge from the topics and expressions preceding them (Lindlof & Taylor, 2002).

The analysis of focus group data presents both challenges and opportunities when compared to other types of qualitative data. While focus groups data can be analyzed in the same manner as interview data (Harding, 2013), there are also unique features of focus groups to be taken into consideration – particularly the opportunity that it provides to observe interactions between group members. Data analysis can take place at the level of the individual or the group.

A fundamental limitation in focus groups (and other forms of qualitative research) is the problem of observer dependency: the results obtained are influenced by the researcher or his or her own reading of the group's discussion, raising questions of validity, called *experimenter's bias*

(Frankfort-Nachmias & Nachmias, 1992). On the other hand, focus groups can create major issues of external validity (Campbell & Stanley, 1963). Other common (and related) criticism involve *groupthink*, where members try to minimize conflict and reach a consensus decision (Irving, 1972), and *social desirability bias*, where members respond in a manner that would be viewed favorably by others (Fisher, 1993).

3.3. RESEARCH STRATEGY

Design science methodology is formed as a chain of six activities that allows researchers to start from any of them and reinitiate the process (Peppers et al., 2007). These activities and their application in this research are described in the table below.

Activities	Application in the research
Problem identification and motivation	The problem of this research is the definition of the impact of Lamarckian evolution in ALife on various scientific disciplines, our life and society. Special attention has been given to the future role of Lamarckian knowledge inheritance in computer science, biology and philosophy
Definition of the objectives	The objectives here are qualitative and aim to identify possible implications of knowledge inheritance across disciplines through Lamarckian evolutionary scheme in artificial organisms
Design and development	The model has been designed using theoretical knowledge about neural networks and genetic algorithms. The model has led to a set of questions, which have been organized and presented to a focus group
Demonstration	The model has been shown to the invited members of the focus group. We have discussed whether this model can be used to simulate Lamarckian evolution in a computational environment and what philosophical significance this model can have, given the possibility of artificial life creation (at this stage, we assumed that the implementation of this model is doable, or at least does not suffer enough roadblocks to preclude attempting it, in order to examine the ethics of pursuing the project)
Evaluation	Using a focus group of researchers from computer science, biology and philosophy, we have validated the model and defined its implications in each mentioned discipline, given the possibility of its implementation in ALife in the near future
Communication	Late-breaking abstract of this research has been submitted and accepted on the scientific conference ALIFE XV. Accepted late-breaking abstracts have been compiled into a Late-Breaking Abstract Proceedings PDF and made publicly available on the conference website. The full article is to be published in IES-2016 conference proceedings.

4. RESEARCH

4.1. FUNDAMENTALS

In 2010, two American biologists Craig Venter and Hamilton Smith have made a bacterium that has an artificial genome—creating a living creature with no ancestor (Gibson et al., 2010). According to Craig Venter himself, this cell has not yet found any practical applications, but it enables a change in philosophy, it is a proof of concept. But the proof of concept that we can potentially create and modify living creatures the way we want was key, otherwise it is just speculation and science fiction.

At the same time, there have been several successful attempts to model Lamarckian evolution in computer science and engineering (Morris et al., 1998; Sasaki and Tokoro, 2000), e.g. in automated docking (Morris et al., 1998). This model of evolution has been used for boosting search in particular kind of applications, however, the cost associated with the evaluation of the objective function with the use of Lamarckian evolution was an issue to consider. Such models have a broad range of applications over several different domains, e.g. optimization in engineering.

If the two mentioned approaches were combined, and there were indeed purely artificial organisms that would learn through interactions with the environment, numerous ethical and philosophical implications would arise. While some may regard the creation of purely artificial organisms as a defining moment in the history of biology, others may claim that the risks could outweigh the benefits.

The inception on species capable of Lamarckian learning does have the potential to do great harm, as well as good. From the philosophical point of view, this research may be seen as playing God and even distorting the essence of life, instead of allowing life to emerge through natural processes and perhaps by nature's will. From a more practical point of view, some irreversible horrors may come creeping out of the flask on the laboratory bench, once a new Lamarckian specimen is introduced to the natural environment. These issues should give pause even to those who normally embrace advances in science with enthusiasm.

4.2. PROPOSAL (MODEL)

Instead of presenting an automaton with single-layer neural networks like in AntFarm (Collins and Jefferson, 1992), we will train a multilayer neural network using backpropagation (of errors) or the generalized delta rule, since a multilayer net can learn any continuous mapping to an arbitrary accuracy (Fausett, 1994). Training a network will include the feedforward of the input training pattern, the backpropagation of the associated error, and the adjustment of the weights.

It is important to note that we used a multilayer neural network with descendant weight updates, as well as backpropagation and delta rule, just to provide an example in an abstract model. There could have been other examples, where other types of neural networks would be used.

Consider the following network (fig. 5), in which we can formulate both feedforward propagation and backpropagation as a series of matrix multiplies. From now on, we are going to index matrices as $A^{(i)}$, where A refers to the type of matrix and (i) is an index of the position of the

matrix in the network (we can also have $(i \rightarrow j)$ for a weight matrix connected layer i to layer j). The only exceptions are the input data matrix X and the output of the network Y . We denote the value of an element in row i and column j of some matrix $A^{(k)}$ with $A_{ij}^{(k)}$ (Dolhansky, 2014).

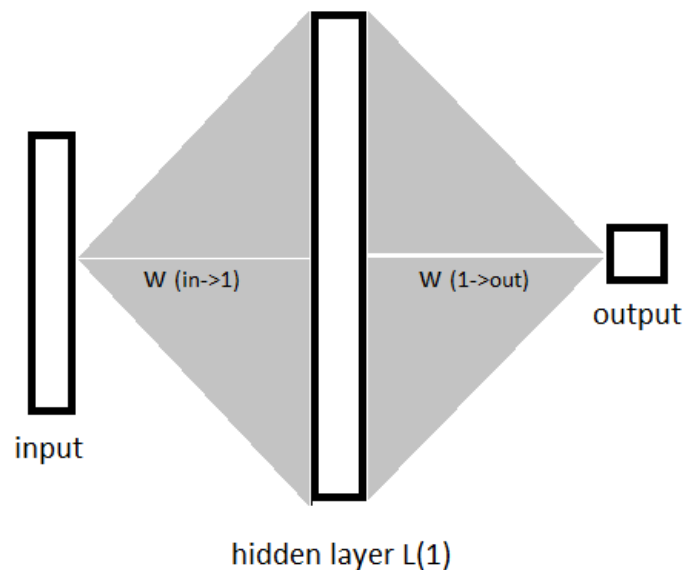


Figure 5: Simplified model of the neural network of an automaton

The defined automata can be made of several areas, part of them would correspond to a neural net (Domeniconi, 1996) and making up a neurological system: vision, hearing, touch, and internal sensing; the rest are physical characteristics. In AntFarm (Collins and Jefferson, 1992), agents of the same colony have identical genetic codes. This is not what we are looking for. The automata in this model have some, although minor differences in their connections between neurons, which better represents such in real biological systems.

The neurological system of automata consists of organs (vision, hearing, touch, internal sensing), through which it receives inputs (like synapses) in the neurons n_m about the environment. Based on the weights $w_{p,q}$, which stand for knowledge, a mathematical function would determine the activation of the neuron. Another function (which may be identical) computes the output of the artificial neuron and lead to an action of automata, like movement, eating, reproducing, or breathing.

Changing weights in the model represents learning. Automaton will learn throughout its life, which means all its weights $w_{i,j}$ will change in order to let it make better decisions (as a result of its learning experience). $C_{w_{i,j}}$ is a learning matrix with weights that change with time. The backpropagation algorithm will be used to compute the necessary corrections. The algorithm can be decomposed in the following four steps (Rojas, 1996): feed-forward computation, backpropagation to the output layer, backpropagation to the hidden layer, and weight updates.

During feedforward, each input unit X_i receives an input signal and broadcasts this signal to the each of the hidden units Z_1, \dots, Z_p . Each hidden unit then computes its activation and sends its signal z_j to output units. Output units Y_k compute their activation y_k to form the response of the net

for the given input pattern (Fausett, 1994). Note that the network will not be fully connected, just like our human brain.

The first step to take will be preparing a population. At the beginning, the weights and thresholds for each individual will be randomized using Nguyen-Widrow algorithm (1990):

1. Define scale factor (β), where n = number of input units, p = number of hidden units:

$$\beta = 0.7(p)^{1/n}$$

2. For each hidden unit ($j = 1, \dots, p$):

- 2.1. Initialize its weight vector by randomizing each weight between -0.5 and 0.5;

- 2.2. Calculate vector length for weights

$$||V_i|| = \sqrt{V_{i1}^2 + V_{i2}^2 + V_{i3}^2 + \dots + V_{in}^2}$$

- 2.3. Update new weights:

$$V_{i \text{ new}} = \frac{\beta \cdot V_{i \text{ old}}}{||V_i||}$$

- 2.4. Calculate threshold values (set biases):

$$\text{Threshold} = \text{random}(-\beta, \beta)$$

The environment is represented as a bi-dimensional plane 500 x 1000, part of the plane is represented on fig. 5 (only part for better visibility), where the automata can move right, left, up and down. The letter A stands for the automata, while O stands for food and Z stands for predators.




Figure 6: Artificial life environment of automata (piece)

Let us provide some exemplary values for better representation. These values can be changed as variables in the parameters section of the simulation software. At the start, we have five predators, four automata and nine food units. The life duration of the automata is 100 time units (t), of a predator – 60t, and food regrows on the same spot in 50t. Automata are born with 50 energy units (E), predators – with 70E, and they spend energy on various actions in a struggle for survival. When the level of energy of A or Z drops to 0, they die.

Every time automata eat food O, their energy increases by 50E, each move decreases it by 10E, reproduction decreases it by 50E, and when A meets a predator Z, A dies. Predators have 70 energy units at birth and they do not consume Os, but if they eat A, their energy increases by 100E.

Food cannot move, predators and automata can move right, left, up and down. The automaton has to learn to come closer to the Os to get energy from them, while moving away from Zs that represent danger. Each time one automaton meets another, they reproduce, and two new automata appear. When Z meets another Z, they reproduce the third Z and lose 30E.

The example of such learning matrix of A_1 that is responsible for vision at the initial time $t=0$ is the following:

$$Vw_{i,j} (t=0) = \{x; y; z; v; u; \dots\}$$

As time goes by, A_1 learns to choose better strategies, and weights change accordingly. This is the example of how a matrix might look like at $t=10$ when some weights have changed:

$$Vw_{i,j} (t=10) = \{x; p; q; v; u; \dots\}$$

At some point, e.g. $t=50$, A_1 and A_2 meet, and in our environment, they have to reproduce and make two new automata. The learning matrix now looks like this:

$$Vw_{i,j} (t=50) = \{p; q; a; b; u; \dots\}$$

Now it is time to see what kind of traits the new life would inherit, according to Darwinian and Lamarckian approaches. In the figure 3, there is an example of what neural networks of parent

A₁ could be. The inputs from the environment reach input neurons on the left, get to a hidden layer, and based on knowledge hidden in weights, output neurons trigger actions on the right.

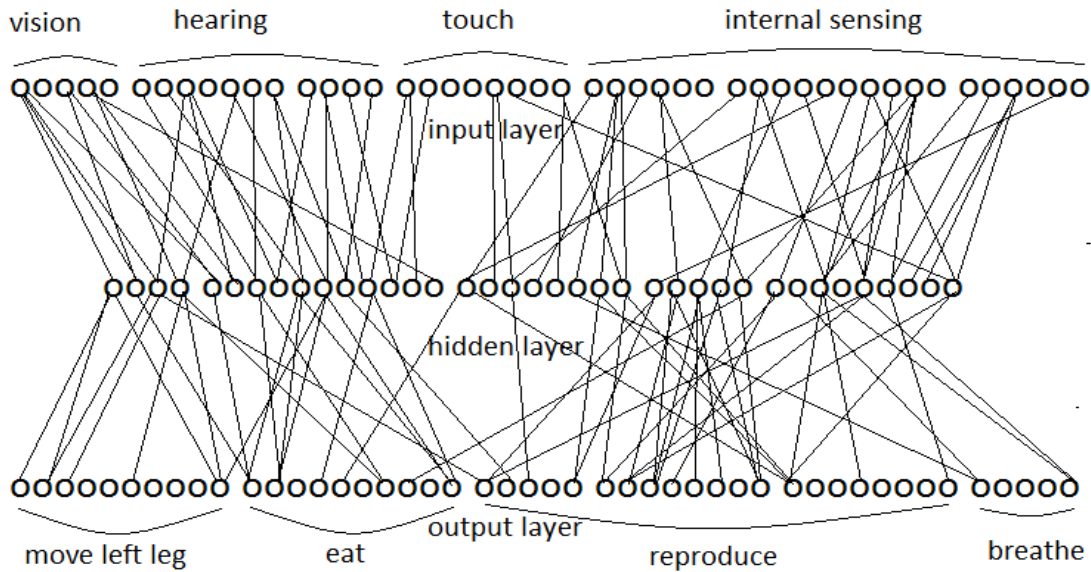


Figure 7: Neural network of parent A₁

The neural network of the second parent would be very similar to such of the first parent. We assume that the amount of the neurons is roughly the same, the difference is in the location of the connections between them, and it should not be very big.

Parents																					
A1	n1	n14	∅	n2	n18	∅	n5	n14	∅	n6	n15	∅	n8	n18	∅	n10	n23	∅	n13	n21	∅
A2	n2	n18	∅	n4	n14	∅	n5	n14	∅	n6	n15	∅	n7	n17	∅	n9	n20	∅	n13	n21	∅
point of crossover																					
Children																					
A3	n1	n14	∅	n2	n18	∅	n5	n14	∅	n6	n15	∅	n7	n17	∅	n9	n20	∅	n13	n21	∅
A4	n2	n18	∅	n4	n14	∅	n5	n14	∅	n6	n15	∅	n8	n18	∅	n10	n23	∅	n13	n21	∅

Figure 8: Darwinian crossover (piece)

According to Darwinist approach, automaton is born with heuristic values that are inherited through a crossover, however, it has almost no knowledge at all (see figure 5). The weights of its neural network are chosen pseudo-randomly. This means that instead of a crossover, the initial weights would be replaced by some heuristic (inborn) values appropriate just for automaton's survival, e.g. breathing or eating. In the figures 8 and 9, there will be only pieces of the whole network in such form (in reality, there are many more neurons, this is the demonstration of the concept).

In a Lamarckian neural network, however, weights will be inherited as well. They can be transferred to children directly (as in figure 9) or through any mathematical function, e.g. average of corresponding weights.

Parents																					
A1	n1	n14	w1,14	n2	n18	w2,18	n5	n14	w5,14	n6	n15	w6,15	n8	n18	w8,18	n10	n23	w10,23	n13	n21	w13,21
A2	n2	n18	w2,18	n4	n14	w4,14	n5	n14	w5,14	n6	n15	w6,15	n7	n17	w7,17	n9	n20	w9,20	n13	n21	w13,21
				point of crossover																	
Children																					
A3	n1	n14	w1,14	n2	n18	w2,18	n5	n14	w5,14	n6	n15	w6,15	n7	n17	w7,17	n9	n20	w9,20	n13	n21	w13,21
A4	n2	n18	w2,18	n4	n14	w4,14	n5	n14	w5,14	n6	n15	w6,15	n8	n18	w8,18	n10	n23	w10,23	n13	n21	w13,21

Figure 9: Lamarckian crossover (piece)

This is just an example to prove an evolutionary neural network combined with Lamarckian approach is possible to create. This model is not attempting to explain how a brain of a pure Lamarckian being would work. It aims to demonstrate a possibility of creating a robot, whose learning would have a Lamarckian component, since at least some of its knowledge would be passed to the offspring.

4.3. PRELIMINARY CONCLUSIONS

In our culture, there is a widespread image of a superior artificial mind conquering the planet. In reality, Darwinian or Lamarckian, an artificial species will more likely become another species’ lunch. The risks about creating artificial life are exaggerated.

When the news about Craig Venter’s achievement came out, people started to become worried about the dangers that ALife could bring. Such worries are of the same nature as popular beliefs that natural is good and artificial is bad. In fact, nothing could be further from the truth – malaria is very natural yet disastrous, while antibiotics are man-made, but very handy sometimes.

However, Lamarckian artificial life can be risky if released into natural environment, given the examples of already existing failed interventions in natural design, mostly because the dimensions of its danger are never known in advance. Therefore, even if the implementation of synthetic Lamarckian evolution is allowed, it should be regulated and licensed in order to avoid malevolent use.

For human beings, however, the possibility of inheriting knowledge would have its major drawbacks. Often our parents are not the people we want to inherit knowledge from, especially in a form of a random combination of their skills instead of a catalog. Moreover, it is nice to inherit knowledge from a person like Einstein, but all the pains, fears and mental issues of our parents would be inherited too, which would have a major negative impact on our lives. In addition, our brains do not work like hard drives that store information the way it is over the years. Our memories get

distorted over time, we forget some details and come up with new ones. Passing such flawed knowledge through generations does not seem like a good idea.

While Lamarckian evolution has wide application in ALife and can be surely used in experiments in artificial environment, for human beings, Darwinian evolution seems to make more sense. At the end of the day, we are all living in dynamic environments where inheriting knowledge is often useless – we keep re-learning again and again, and the ability to unlearn old skills and study everything anew seems to be the new literacy.

All these hypotheses and assumptions have led to many questions that we organized and presented to our focus group. The discussion will provide better understandings whether these concerns are valid.

4.4. VALIDATION

Our focus group discussion constituted gathering together experts from different backgrounds and experiences to discuss the topic of interest, which is the introduction of Lamarckian species to the world. The group of participants has been guided by a moderator who introduced questions for discussion and helped the group to participate in a lively and natural discussion amongst themselves.

The validity of the discussion relies on allowing the participants to agree or disagree with each other so that it provides an insight into how a group thinks about an issue, about the range of opinion and ideas, and the inconsistencies and variation that exists in a particular community in terms of beliefs and their experiences and practices (Viji and Benedict, 2014).

Compared to surveys, focus groups can reveal a wealth of more detailed information and deeper insight. If executed right, a focus group creates a comfortable environment that puts participants at ease, inviting them to consciously answer questions in their own words and add personal meaning to their answers.

The size of the group has been chosen in a way that the group is big enough to generate rich discussion, but not so big that some participants are left out. We invited two philosophers, two researchers from computer science, and one biologist. The participants have been chosen from different backgrounds to give a better perspective in an open discussion, rather than just a simple record of their personal attitudes to the question. The goal was to generate a maximum number of different ideas and opinions in the defined time (Eliot & Associates, 2005).

The questions have been grouped by discipline and presented to the focus group in the following format:

Discipline	Questions
Biology	<ol style="list-style-type: none"> 1. Would the introduction of Lamarckian organisms become a determining moment in the history of biology? 2. Can this research be abused to create a biological weapon? 3. Would the benefits of Lamarckian evolution be significant in synthetic biology? 4. If inheriting knowledge was possible in the same way as we inherit eye or hair color, does it mean that we would inherit fears and psychological disorders from our parents too? 5. Do you think we would be better fit to our environment if we could inherit knowledge from parents? 6. How do you see the impact on evolution if Lamarckian knowledge inheritance was real? What would happen to the rate of such evolution, would it speed up or slow down? Why? 7. What potential risks can be contained?
Computer science	<ol style="list-style-type: none"> 1. What kind of systems can be developed using such evolution? When, now, in 10 years? 2. Will this evolution have any practical application in engineering in the nearest future? Why? 3. Do you agree that adding weights to the model is the correct way to model Lamarckian inheritance? Is anything missing?
Philosophy	<ol style="list-style-type: none"> 1. If we are building artificial life models, are we playing the role of gods? 2. If we are able to create a different life using such type of models, do we have a chance to be better gods? 3. Is creating artificial life consistent with nature/God? 4. Does this evolution contradict the idea of free will/freedom of choice? 5. Do we have a chance to be happier if we could know what our parents know? 6. Is Lamarckian evolution more ethical than Darwinian? Why? 7. Will this type of research help us understand the relationship between us and the creator any better, or this is just a cybernetic adventure?

5. DISCUSSION

Focus groups are conducted based on a set of structured and predefined questions, but the discussion is free-flowing. Ideally, participant comments will stimulate and influence the thinking and sharing of others. Some people might even change their thoughts and opinions during the group discussion.

In order to make use of all participant comments, it is essential to boil them down to essential information using a systematic and verifiable process. In our case, we have started by transcribing all focus group tapes and inserting notes into transcribed material where appropriate. Then the transcripts have been cleaned up by stripping off nonessential words (Eliot & Associates, 2005).

The answers presented below summarize all responses and disagreements that have been recorded during the focus group discussion. We noted that the researchers from philosophy sometimes disagreed with each other, hence their opinions are represented below as A and B. In fact, we found their opinions complementary rather than contradictory.

5.1. BIOLOGICAL IMPACT

1. Would the introduction of Lamarckian organisms become a determining moment in the history of biology?

Humans engineered organisms for centuries; synthetic life is just one step further, which makes it, obviously, a great achievement, yet this can be hardly considered a determining moment. We already have powerful means to engineer organisms, and at this moment of time, synthetic life does not add that much. Moreover, Craig Venter and his team did not actually engineer synthetic life, the resulting cell only has synthetic DNA that requires a living host cell.

2. Can this research be abused to create a biological weapon?

At this very moment, nobody knows for sure where it will all lead. Like in any type of research, some of it can be abused by malevolent scientists, but that has been true of just about every human advance and should not stop us. Any scientific advance can be abused; Lamarckian evolution in synthetic life no different.

3. Would the benefits of Lamarckian evolution be significant in synthetic biology?

The benefits are rather unclear, but they look promising. The approach of designing synthetic cells will allow us to begin with a DNA sequence and engineer organisms precisely up to our specific needs, including building Lamarckian species. At the moment, we can think of better drugs, greener fuels and dramatic changes in chemical industry, and Lamarckian learning can potentially benefit the research in this direction.

4. If inheriting knowledge was possible in the same way as we inherit eye or hair color, does it mean that we would inherit fears and psychological disorders from our parents too?

Yes, but in the same way that we inherit any other disease or trait. For example, we inherited fear of snakes from our ancestors, which is positive, since snakes can be deadly for us. This does not seem to be a particularly dangerous thing.

5. Do you think we would be better fit to our environment if we could inherit knowledge from parents?

Not really. In fact, it is questionable whether artificial Lamarckian organisms can survive in natural environment at all. There is no evidence that knowledge inheritance to such degree can enable Lamarckian species to prosper in a constantly changing environment like real world.

6. How do you see the impact on evolution if Lamarckian knowledge inheritance was real? What would happen to the rate of such evolution, would it speed up or slow down? Why?

It depends on what knowledge is being inherited. The knowledge that is necessary for the survival of the species is being already inherited – as it was already mentioned, organisms are not born blank. It is unclear how all extra knowledge would benefit the species. At least human brain, despite a common belief, is nothing like a hard drive of the computer – one piece of knowledge does not occupy the space where some other knowledge could be stored, yet there is no evidence that learning from scratch would be slower or faster than relearning.

7. What potential risks can be contained?

On the one hand, synthetic organisms carry no greater risks than natural ones, even though new species that evolve via Darwinian evolution do not receive nearly as much recognition or cause fear as hybridization by genetic engineering. On the other hand, there are some terrifying examples of other types of artificial constructions, like GEO crops and over-use of pesticides that led to environmental problems. In addition, the research in synthetic life is expensive and may be driven by profits more than benefits to humanity.

5.2. TECHNICAL IMPACT

1. What kind of systems can be developed using such evolution? When, now, in 10 years?

Lamarckian evolutionary algorithms have been used for quite some years, and quite successfully. However, their application in strong ALife or artificially engineered organisms we would consider *alive*, can only be approximated only based on the existing roadmaps of futurists. These assumptions would have little precision, and the feasibility of this implementation will depend on the accessibility of brain-emulating hardware, which is a hard thing to do at sufficiently low cost.

2. Will this evolution have any practical application in engineering in the nearest future? Why?

As mentioned before, it already does, and its potential is quite huge. Since 80s, it has been used for function optimization; this evolutionary algorithm has proven to be robust and fit for various optimization problems and machine learning techniques.

3. Do you agree that adding weights to the model is the correct way to model Lamarckian inheritance? Is anything missing?

We believe this is one of the simplistic ways of representing Lamarckian evolution, and alternative models are possible also. At the same time, for now this research is only related to computer science,

hence, it is rather a cybernetic adventure, not a way to understand the relationship between us and the creator any better.

5.3. PHILOSOPHICAL IMPACT

1. If we are building artificial life models, are we playing the role of gods?

A: If we are building ALife models, we cannot say that we are playing the role of gods. Can we define God? We do not know what God is. If we do not know, how can we compare?

B: Maybe we are playing the role of gods when we are creating such artificial life models, but only if we understand God as a metaphor. We imitate gods in the same creationist sense as artists do when they write a musical or literary composition.

2. If we are able to create a different life using such type of models, do we have a chance to be better gods?

A: Who knows? Without the proper definition of God, it is impossible to answer this question.

B: If we are to take already existing definitions of God, we would not necessarily be better gods, at least according to Christianity, because the God in a Christian sense is already a perfection – kind and good, while we, humans, are not perfect. If we define God as Stephen Hawking does, which means it has the knowledge of the past, present and future, it would be difficult to compete with that. In terms of better capacity, however, artificial life has certainly better potential than human race, because it can be faster and more efficient by a range of indicators.

3. Is creating artificial life consistent with nature/God?

B: Absolutely. Humanity has been playing the role of God in creationist sense throughout the history, and there is no reason to limit nature's creativity to exclude acts of man. Since humans are creatures of God, their actions are part of God's or nature's design. If nature has given humans brains strong enough that they will invent species that supersede or eliminate them one day, this can be seen as part of evolution, or God's plan, whatever you may call it.

4. Does this evolution contradict the idea of free will/freedom of choice?

B: If we inherit knowledge from parents, the concept of free will may disappear. By definition, free will is the capacity for an agent to make choices in which the outcome has not been determined by past events. The model of knowledge inheritance contradicts the fundamental idea of free will. In addition, we would not have to learn many things, since the answer would already be in our heads, therefore, it seems like there would be less trial and error.

A: On the other hand, this evolution does not necessarily contradict free will or freedom of choice in the practical sense. What matters is not if you have more knowledge, but what you do with this knowledge. In addition, according to Plato, we do not learn anything, we simply recall, so is there any free will at all?

5. Do we have a chance to be happier if we could know what our parents know?

B: Such evolution also does not guarantee that we would be happier if we could inherit knowledge. The joy of knowledge discovery would be taken away.

A: Imagine you are 5 years old and you have the knowledge of a 25-year-old person, how will you react when you realize the limitations of your body and disproportional development? A disembodied mind will quickly get depressed! If everything develops proportionally, evolution will be faster and nothing more. We need time to process all the knowledge as well. Happiness will not be impacted. Also, even though Western philosophy tends to imply that possessing more knowledge will lead us to happier lives, in practice, it does not always work this way. Besides knowledge, happiness includes factor of pleasure, and sometimes that means that ignorance is bliss. Moreover, by definition, knowledge lies in a context. If you inherit knowledge, not just data/information, you will have a double personality inside yourself. Do you want it?

6. Is Lamarckian evolution more ethical than Darwinian? Why?

B: If ethics implies maximization of well-being of all conscious organisms, and Lamarckian evolution can achieve that by inheriting knowledge, this may be true, but does not seem obvious. From the first sight, Lamarckian evolution looks very convenient. In fact, our cultural evolution is already Lamarckian, as we can culturally pass what we learned to our offspring. If we take a look at our society several centuries ago, we can notice how much we have evolved due to knowledge transferred from one generation to another, so it seems like rather an ethical thing.

7. Will this type of research help us understand the relationship between us and the creator any better, or this is just a cybernetic adventure?

B: Artificial life is limited to a rational sense of the human being, does not have the spiritual dimension of the unknown, and cannot embrace the concept of a God, while humans even with underdeveloped reasoning skills (like children) can open themselves to irrational things like religion. No matter how our rational capabilities increase, they will not let us achieve gods or universe. The access to God is only possible through mysticism, which is a different dimension of a human being that is not possible to create in AI.

5.4. ANALYZING FOCUS GROUP DATA

The main objective of our focus group data analysis is to fairly represent the data and communicate what the data reveal given the purpose of the study. However, focus group analysis is different from quantitative analysis in numerous ways, and even compared to other qualitative analysis strategies, focus groups have their own peculiarities. While interviews and speeches are usually logically structured, during a focus group discussion, the moderator will likely hear spontaneous and inconsistent comments, people changing their minds, wandering conversations and emotionally intense answers that may influence other participants' behavior.

While evaluating biological and technical impact comes down to finding relevant references with the help of the researchers' experience in the field, philosophical argumentation is a fundamentally different task, based on different kind of methods.

When a researcher from philosophy answers the question about a possibility of a human to be a better God, it is important to remember that deity is a concept conceived in diverse ways in various cultures, typically as a natural or supernatural being considered divine or sacred (O'Brien, 2009). Nevertheless, the views on the definition of God differ tremendously depending on a person being interviewed, thus the response may lack objectivity.

The whole idea of a consistency with nature or God's will is also just an idea. We cannot say for sure if there is a creator and what was his or her idea or intention. In addition, the assumption about God as the creator of everything generates an age-old problem: it is not coherent to argue that the universe was created by God, but God was in turn created by God to the second power, who was in turn created by God to the third power, and so on.

The sciences have grown steadily bolder in their claim that all human behavior can be explained through the clockwork laws of cause and effect. This shift in perception is the extension of an intellectual revolution that began about the time when Charles Darwin first published "On the Origin of Species". Soon after Darwin offered his theory of evolution, his cousin Sir Francis Galton began to draw out the implications: if species have evolved, then mental abilities like intelligence must be hereditary. But we use these abilities—which some people possess to a greater degree than others—in decision-making. So our ability to choose our destiny is not free, but depends on our biological inheritance (Cave, 2016), which only reinforces the idea of determinism and shatters the concept of free will. If the concept of free will disappears along with scientific advancements, we seem to have little reason to worry about the influence of implementation of Lamarckian learning on this concept.

The aspect of happiness in Lamarckian knowledge inheritance is also questionable. One of the focus group participants stated that the joy of discovering knowledge will be taken away, which does not seem to be true, as there is plenty more knowledge to discover apart from what our parents may know. Moreover, happiness is not equal to joy and may have other components such as meaningfulness or having a purpose, and knowledge inheritance, thus probably better awareness, might actually help in that matter.

At the end, the statement that artificial life cannot have spiritual experiences and is limited to a rational sense is also just an opinion. The supporters of strong AI may strongly disagree with this. On the other hand, we are not even sure whether spiritual experiences are possible in humans – those might be just hallucinations caused by external or internal influences, and a bug in the system may cause similar effect in AI.

6. CONCLUSIONS AND FUTURE WORK

6.1. MAIN CONTRIBUTIONS

This research constitutes a position paper that presents an opinion about future application of Lamarckian learning in various disciplines based on existing research in the field and expert opinions shared in a focus group discussion. This research aims to demonstrate a specific point of view on a future matter that was formed during this discussion, and the purpose of this dissertation is to defend, explain and document the reasoning behind that position. The analysis of the discussion has led us to several conclusions that are represented below.

It seems that Lamarckian organisms can have wide practical application across several different domains, therefore this type of research should rather be allowed and encouraged. But even though Lamarckian evolutionary algorithm already holds major benefits for humanity and promises even more, this implementation needs regulation. For now, potential benefits seem to outweigh risks, however, the risks are unknown and the required investment might be an issue.

Firstly, it is absolutely necessary to prevent malevolent use of the research. It can be abused in numerous ways, e.g. applied in the creation of biological weapon and bioterrorism. Secondly, since the project will require a lot of investment, the research may become driven by profits for harmful purposes rather than benefits to humanity as a whole. Thirdly, safety measures should be taken before releasing the Lamarckian species into natural environment. This will be needed not only for the sake of safety of natural biodiversity, but also to help the Lamarckian organisms survive outside the lab.

Nowadays, even for a non-scientist it seems quite easy to distinguish a living organism from a non-living, except probably for viruses, whose status is still questionable. The agreed on definition of life, however, does not exist, and the inception of artificial species, in this case, artificial Lamarckian species makes it even murkier. It is still unclear whether this implementation will defeat the divinity of life or concept of the soul and prove that there is no magic spirit of vitality. Until now, man did not manage to create life from scratch, only to manipulate it, so the question remains open.

The inception of Lamarckian organisms may eliminate some existing philosophical concepts, such as free will, because inherited knowledge contradicts its definition, however, what matters is the *ability* to choose which stays put. In this sense, the definition of free will in philosophy might be changed or expanded, considering the possibilities that are arising in artificial life.

This type of research is still rather related to computer science and does not prove or disprove the theory of creationism. Therefore, by conducting this kind of studies, we may not embrace the concept of God or understand the origins of our species. Also we may not understand the relationship between us and the creator any better, as well as find out whether the creator has ever existed at all.

6.2. LIMITATIONS OF THE CURRENT WORK

A homogeneous group of strangers is supposed to comprise a focus group. Homogeneity levels the playing field and reduces inhibitions among people who will probably never see each other again (Eliot & Associates, 2005). Our group has been not entirely homogeneous. The researchers were from different fields, and better results would be obtained in a homogeneous group that consists entirely of biologists or entirely of philosophers. Also, best results are expected in a group where participants do not know each other (Eliot & Associates, 2005), which was not the case with participants from computer science.

Moreover, it takes more than one focus group on any one topic to produce valid results – usually three or four. What can demonstrate that enough focus group have been conducted (with the same set of questions) is when a moderator does not hear anything new anymore, i.e. reaches a point of saturation (Eliot & Associates, 2005). In our case, only one focus group was organized, and it surely had its impact on the precision of the results.

This focus group has provided some good insights, yet this method has some drawbacks. Firstly, some participants were not fully prepared for discussing this topic, because the topic itself is rather complex and requires at least some level of expertise in the field or closely related fields. Secondly, some opinions were outdated (e.g. the statement that artificial life is limited to a rational sense of the human being and cannot have spiritual experiences) and the concepts of AI and ALife were used as if it was the same thing. Moreover, when it comes to philosophy, without a proper preliminary research, the answers of the participants may lack objectivity and can be highly affected by personal beliefs (e.g. focusing on the Christian definition of God and not considering any other).

6.3. FUTURE WORK

Ideally, the qualitative analysis in focus group should be continuous; it only begins in the first focus group, and future work will be needed for better precision. The results from the first focus group may suggest topics to emphasize or include in later focus groups. For example, based on the initial focus group, it is important to determine if central questions are still relevant. In our research, we have found that Lamarckian learning is already widely applied in engineering, therefore the question whether such implementation is possible has become obsolete. Also such assumptions as Lamarckian organisms taking over the world and threatening the well-being of humanity still remain fantasies and are too unlikely to happen. What is important for us to know, and what we are going to focus our next research on, is the future application of Lamarckian systems in our daily life. Therefore, our questionnaire will be changed accordingly.

In future research, not only we would expand the size of the focus group and conduct more groups, but choose more informed and prepared participants, probably using some incentives. Also we would like to organize more homogeneous groups instead, such that the participants would be from the same field, with the same academic status, and would not know one another (to ensure maximum disclosure and better representation of valid opinions).

While this dissertation has demonstrated the potential of efficient application of Lamarckian learning algorithm, many opportunities for extending the scope of this thesis remain. This section presents some of these directions.

In the last two decades, we have seen the examples of usage of the Lamarckian learning algorithm in “soft” artificial life. This research, however, was related exclusively to computer science, and the impact of training intelligent agents using Lamarckian learning did not extend beyond the artificial environment where it operated. Considering the advances in aforementioned problems, Lamarckian learning can be implemented in “hard” and “wet” artificial organisms. This means that it can lead to a creation of a physical self-reproducing robot capable of Lamarckian knowledge transmission, or a Lamarckian biological organism with artificially engineered DNA (to date, there is no evidence that it can be done without using a biological host cell).

Recent work in self-replicating systems research includes self-replicating rapid prototypers (Freitas et al., 2010), NASA Institute for Advanced Concepts studies on accelerating space exploration (Hod & Malone, 2007) and architecture of unmanned lunar factories (Chirikjian et al., 2002), New York University artificial DNA tile motifs (Wang et al., 2011) and others. Implementation of Lamarckian learning may improve performance of these systems.

What we can expect from better understanding of life is improved decision making on all levels: managing ecological resources, regulating social interactions, planning urban systems, commercializing biotechnology, and more (Aguilar et al., 2014).

7. BIBLIOGRAPHY

- Aguilar, W., Santamaria-a-Bonfil, G., Froese, T., & Gershenson, C. (2014). The Past, Present, and Future of Artificial Life. *Frontiers in Robotics and AI*, 1.
- Ayala, F. J. (2008). Science, evolution, and creationism. *Proceedings of the National Academy of Sciences of the United States of America* (Vol. 105).
- Bar-Yam, Y. (1997). *Dynamics of Complex Systems*. Reading, MA: Addison-Wesley.
- Bedau, M. A. (2002). Leonardo. *The Scientific and Philosophical Scope of Artificial Life*. MIT Press Vol. 35, No. 4, p. 395-400.
- Bedau, M. A. (2003). Artificial life: organization, adaptation and complexity from the bottom up. *Trends Cogn. Sci. (Regul. Ed.)* 7, 505–512.
- Bedau, M. A., McCaskill, J. S., Packard, N. H., and Rasmussen, S. (2000). Open Problems in Artificial Life. *Massachusetts Institute of Technology. Artificial Life* 6: p. 363–376.
- Borresen, J. and Lynch, S. (2009). Neuronal computers. *Nonlinear Analysis, Theory, Methods and Applications*, 71(12).
- Bourne, P. E., Brenner, S. E., and Eisen, M. B. (2005). PLoS computational biology: a new community journal. *PLoS Comput. Biol.* 1:e4.
- Bray, D. (2011). *Wetware: A Computer in Every Living Cell*. Yale University Press: New Haven, Connecticut.
- Bull, J. J. and Wichman, H. A. (2001). Applied Evolution. *Annual Review of Ecology, Evolution, and Systematics*, 32, 183–217.
- Campbell, D. T., & Stanley, J. C. (1963). *Experimental and Quasi-Experimental Designs for Research*. Houghton Mifflin Company.
- Cave, S. (2016). There's No Such Thing as Free Will. *The Atlantic*, 06(16).
- Chirikjian, G. S., Zhou, Y., & Suthakorn, J. (2002). Self-replicating robots for lunar development. *IEEE/ASME Transactions on Mechatronics*, 7(4), 462–472.
- Collins, R. J. and Jefferson, D. (1992). Antfarm: towards simulated evolution. In *Artificial Life II* (pp. 579–601). Retrieved from <http://citeseer.ist.psu.edu/collins91antfarm.html>
- Cortez, P., Rocha, M., and Neves, J. (2002). A Lamarckian approach for neural network training. *Neural Processing Letters*, 15(2), 105–116.
- Darwin, C. (1859). *On the Origin of Species by Means of Natural Selection, or the Preservation of Favoured Races in the Struggle for Life*.
- Davidor, Y. (1991). *Genetic Algorithms and Robotics*. World Scientific Series in Robotics and Intelligent Systems: Volume 1. (Weizmann Inst. Sci., Israel).

- Demuth, H. (2002). Neural Network Toolbox. *Networks*, 24(1), 1–8.
- Diallo, S. Y., Padilla, J. J., Bozkurt, I., and Tolk, A. (2013). Modeling and simulation as a theory building paradigm. In *Ontology, Epistemology, and Teleology for Modeling and Simulation*. Springer, p. 193–206.
- Djurfeldt, M., Lundqvist, M., Johansson, C., Rehn, M., Ekeberg, O., and Lansner, A. (2008). Brain-scale simulation of the neocortex on the IBM Blue Gene/L supercomputer. *IBM Journal of Research and Development*, 52, 31–41.
- Dolhansky, B. (2014). Artificial Neural Networks: Matrix Form (Part 5). *ML Primers, Neural Networks*. Retrieved from <http://briandolhansky.com/blog/2014/10/30/artificial-neural-networks-matrix-form-part-5>
- Domeniconi, C. (1996). Proposal of a Darwin-Neural Network for a Robot Implementation. *Perspectives in Neural Computing*, Taylor Ed., p. 186–193.
- Edelman, G. (1987). Neural Darwinism. *Neural Darwinism (Vol. 4)*.
- Eliasmith, C., Stewart, T. C., Choo, X., Bekolay, T., DeWolf, T., Tang, Y., and Rasmussen, D. (2012). A large-scale model of the functioning brain. *Science*, 338, 1202–1205.
- Eliot & Associates. (2005). Guidelines for conducting a focus group. Available at: http://assessment.aas.duke.edu/documents/How_to_Conduct_a_Focus_Group.pdf.
- Fahmy, M. (2014). Artificial Life and the Philosophy of Science. *ALIFE 14: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*.
- Farmer, J.D. and Belin, A. (1990). *Artificial Life: The Coming Evolution*. Santa Fe Institute. WORKING PAPER.
- Fausett, L. (1994). *Fundamentals of Neural Networks: Architectures, Algorithms, and Applications*. Prentice-Hall.
- Fisher, R. J. (1993). Social Desirability Bias and the Validity of Indirect Questioning. *Journal of Consumer Research*, 20(2), 303.
- Frankfort-Nachmias, C., & Nachmias, D. (1992). *Research Methods in the Social Sciences*. *American Political Science Review*, 25, 600.
- Fraser, A. S. (1958). Monte Carlo analyses of genetic models. *Nature* 181 (4603): 208–9.
- Freitas, V., Queijo, L., & Lima, R. (2010). Rapid prototyping of 3D anatomical models to hemodynamic studies. *International Conference on Biomedical Electronics and Devices*, 246–251.
- Gershenson, C. (2003). *Artificial Neural Networks for Beginners*. Cornell University Library.
- Gibson, D. G., Glass, J. I., Lartigue, C., Noskov, V. N., Chuang, R.-Y., Algire, M. a, ... Venter, J. C. (2010). Creation of a bacterial cell controlled by a chemically synthesized genome. *Science (New York, N.Y.)*, 329(5987), 52–56.

- Grefenstette, J. J., Ramsey, C. L., and Schultz, A. C. (1990). Learning sequential decision rules using simulation models and competition. *Machine Learning*, 5(4), 355–381.
- Hall, B.K. and Hallgrimsson, B. (2007). Strickberger's Evolution. *Molecular Nutrition & Food Research*, 58(4), 84–902.
- Harding, J. (2013). *Qualitative Data Analysis from Start to Finish*. London, SAGE Publishers.
- Hevner, A. R., March, S. T., Park, J., and Ram, S. (2004). Design Science in Information Systems Research. *MIS Quarterly*, 28, 75-106.
- Hewitson, B. C. and Crane, R. C. (1994). *Neural Nets: Application in Geography*. Kluwer Academic Publishers.
- Hinton, G. E. and Nowlan, S. J. (1987). How Learning Can Guide Evolution. *Complex Systems*, 1, 495–502.
- Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems*. Ann Arbor MI University of Michigan Press (Vol. Ann Arbor).
- Huerta, M., Haseltine, F., Liu, Y., Downing, G., and Seto, B. (2000). NIH Working Definition of Bioinformatics and Computational Biology. BISTIC Definition Committee.
- Husbands, P., Harvey, I., Cliff, D., and Miller, G. (1997). Artificial evolution: a new path for artificial intelligence? *Brain and cognition* (Vol. 34).
- Jamshidi, M. (2003). Tools for intelligent control: fuzzy controllers, neural networks and genetic algorithms. *Philosophical Transactions of the Royal Society A* 361 (1809): 1781–808.
- Janis, I. L. (1972). *Victims of Groupthink: Psychological Studies of Policy Decisions and Fiascoes*. Houghton, Mifflin.
- Keeley, B. L. (1998). Artificial life for philosophers. *Philosophical Psychology*, 11(2), 251–260.
- Koza, J. R. (1992). *Genetic Programming*. MIT Press.
- Kumar, D. and Bhatnagar, R. (2010). An approach implements artificial intelligence into human life with new technologies and application. *International Journal on Emerging Technologies* 1(2).
- Kurzweil, R. (1999), *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*. New York, NY: Penguin Books.
- Lamarck, J. B. (1809). *Philosophie zoologique, ou Exposition des considérations relatives à l'histoire naturelle des animaux*. Paris.
- Langton, C. G. (ed.) (1989). *Artificial Life: Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*. Los Alamos: Addison-Wesley. *Complex Adaptive Systems*.
- Langton, C. G. (1998). A new definition of artificial life. Available at: <http://scifunam.fisica.unam.mx/mir/langton.pdf>

Lindlof, T. R., & Taylor, B. C. (2002). *Qualitative Communication Research Methods*, 2nd Edition. Thousand Oaks, CA: Sage.

Lipson, H. & Malone, E. (2007). *Autonomous Self-Extending Machines for Accelerating Space Exploration*. NIAC CP 01-02 Phase I.

Lewis, P.R., Platzner, M., Rinner, B., Torresen, J., and Yao, X. (2016). *Self-aware Computing Systems. An Engineering Approach*. Springer, Natural Computing Series.

Maes, P. (1993). Modeling adaptive autonomous agents. 1, 135–162.

Manderick, B. (1991). *Selectionism as a Basis of Categorization and Adaptive Behavior*. Ph.D. Thesis, AI Lab VUB Brussels.

Marczyk, A. (2004). *Genetic Algorithms and Evolutionary Computation*. The TalkOrigins Archive.

Markram, H. (2006). The blue brain project. *Nature Reviews Neuroscience*, 7, 153–160.

Mitchell, M. (1996). *An Introduction to Genetic Algorithms*. The MIT Press. Cambridge, Massachusetts.

Mitchell, T. M. (1997). *Machine Learning*. McGraw-Hill Science/Engineering/Math.

Morgan, D. L. (1997). *Focus groups as qualitative research / David L. Morgan*. Qualitative research methods series (Vol. 16).

Moritz, C. (2010). *Biology 1B—Evolution Lecture 1. Introduction to Evolution*. UC Berkeley.

Morris, G. M., Goodsell, D. S., Halliday, R. S., Huey, R., Hart, W. E., Belew, R. K., and Olson, A. J. (1998). Automated Docking Using a Lamarckian Genetic Algorithm and an Empirical Binding Free Energy Function. *Journal of Computational Chemistry*, 19, 1639–1662.

Neumann, J. v. (1963). *The General and Logical Theory of Automata*. *Collected Works Volume V: Design of Computers, Theory of Automata and Numerical Analysis*. Ed. A.H. Taub. New York: Pergamon, p. 288–326.

Nguyen, D., & Widrow, B. (1990). Improving the learning speed of 2-layer neural networks by choosing initial values of the adaptive weights. *IJCNN Int. Joint Conf. Neural Networks*, 13, C21.

O'Brien, J. (2009). *Encyclopedia of Gender and Society*. SAGE Publications. p. 191. ISBN 978-1-4129-0916-7.

Olson, E. T. (1997). The ontological basis of strong artificial life. *Artificial Life*, 3(1), p. 29–39.

Pilpel, Y. (2016). *Realizing Lamarckian Evolution*. Center for Bits and Atoms. MIT Media Lab, E14-633.

Preissl, R., Wong, T. M., Datta, P., Flickner, M. D., Singh, R., Esser, S. K., ... Modha, D. S. (2012). *Compass: A scalable simulator for an architecture for cognitive computing*. *Proceedings of Supercomputing 2012, Salt Lake City, November 10 –16, 2012*.

- Rechenberg, I. (1973). *Evolutionsstrategie – Optimierung technischer Systeme nach Prinzipien der biologischen Evolution* (PhD thesis) (in German). Fromman-Holzboog.
- Rojas, R. (1996). *Neural networks: a systematic introduction*. Neural Networks, 502.
- Ross, B. J. (1999). A Lamarckian Evolution Strategy for Genetic Algorithms. In Lance D. Chambers, editor, *Practical Handbook of Genetic Algorithms: Complex Coding Systems*, Boca Raton, FL: CRC Press, volume III, pp. 1-16.
- Rumelhart, D. and McClelland, J. (1986). *Parallel Distributed Processing*. MIT Press, Cambridge, Mass.
- Russell, B. (1959). *The Problems of Philosophy* [1912]. Oxford: Oxford University Press, 11, 14.
- Sandberg, A. (2014). Ethics of brain emulations. *Journal of Experimental & Theoretical Artificial Intelligence*, (May), 1–19.
- Sandberg, A. and Bostrom, N. (2008): *Whole Brain Emulation: A Roadmap Technical Report*. Future of Humanity Institute, Oxford University.
- Sasaki, T. and Tokoro, M. (2000). Comparison between Lamarckian and Darwinian Evolution on a Model Using Neural Networks and GAs. *Knowledge and Information Systems (KAIS): An International Journal*, Springer.
- Sellars, W. (1963). *Philosophy and The Scientific Image of Man*. *Science, perception and reality* 2: p. 35-78.
- Sharma, V. P. (2010). *Nature at Work - the Ongoing Saga of Evolution*. Springer, p. 121
- Simon, H. A. (1996). *The Sciences of the Artificial* (3rd ed.) MIT Press, Cambridge, MA.
- Springer, J. T. and Holley, D. (2013). *An Introduction to Zoology*. Jones & Bartlett Publishers, p.609
- Stergiou, C. and Siganos, D. (1996). *Neural Networks*. Retrieved March 2008.
- Sterling, B. (1992). *Artificial life*. *The Magazine of Fantasy and Science Fiction*.
- Stewart, R. C. (2005). *The Journal of Evolutionary Philosophy*. The Academy of Evolutionary Metaphysics.
- Taylor, C. (1992). *To Follow a Rule*. M. Hjort Ed. *Rules and Conventions – Literature, Philosophy, Social Theory*. Baltimore, Md.: John Hopkins University Press, p.168-185.
- Vaishnavi, V. and Kuechler, W. (2007). *Design Research in Information Systems*. Wwwwisworldorg.
- Van Aken, J. E. (2005). Management research as a design science: Articulating the research products of mode 2 knowledge production in management. *Br J Manage*. 16(1): 19–36.
- Viji, V. and Benedict, K. (2014). A Narrative Interpretation of a Focus Group Discussion Episode on Emerging Educational Taxonomies by a Novice Investigator. *International Journal of Humanities and Social Science Invention* ISSN. 3(7) pp: 2319-7722.

Wang, B. (2010). Ray Kurzweil Responds to the Issue of Accuracy of His Predictions. Next Big Future. Coverage of Disruptive Science and Technology.

Wang, T., Sha, R., Dreyfus, R., Leunissen, M. E., Maass, C., Pine, D. J., ... Seeman, N. C. (2011). Self-replication of information-bearing nanoscale patterns. *Nature*, 478(7368), 225–228.

Weiss, M. L. (2007). *Neuronal Network Research Horizons*. Nova Science Publishers Inc., New York, p.7

Wilkins, J. S. (1997). Evolution and Philosophy: Does evolution make might right? TalkOrigins Archive, 22.

Wilkinson, S. (2004). Focus groups: A feminist method. In S.N. Hesse-Biber & M.L. Yaiser (eds.), *Feminist perspectives on social research*. New York, Oxford University Press.