

ALF – A Fitness-Based Artificial Life Form for Evolving Large-Scale Neural Networks

Rune Krauss Marcel Merten Mirco Bockholt Rolf Drechsler

Group of Computer Architecture, University of Bremen, Germany

Cyber-Physical Systems, DFKI GmbH, Germany

{krauss, mar_mer, bockholt, drechsler}@uni-bremen.de

ABSTRACT

Machine Learning (ML) is becoming increasingly important in daily life. In this context, *Artificial Neural Networks* (ANNs) are a popular approach within ML methods to realize an artificial intelligence. Usually, the topology of ANNs is predetermined. However, there are problems where it is difficult to find a suitable topology. Therefore, *Topology and Weight Evolving Artificial Neural Network* (TWEANN) algorithms have been developed that can find ANN topologies and weights using genetic algorithms. A well-known downside for large-scale problems is that TWEANN algorithms often evolve inefficient ANNs and require long runtimes.

To address this issue, we propose a new TWEANN algorithm called *Artificial Life Form* (ALF) with the following technical advancements: (1) speciation via structural and semantic similarity to form better candidate solutions, (2) dynamic adaptation of the observed candidate solutions for better convergence properties, and (3) integration of solution quality into genetic reproduction to increase the probability of optimization success. Experiments on large-scale ML problems confirm that these approaches allow the fast solving of these problems and lead to efficient evolved ANNs.

1 INTRODUCTION

In daily life, the usage of *Machine Learning* (ML) is becoming more popular, especially *Reinforcement Learning* (RL) [18]. As an important ML method, RL improves the decision-making ability of an intelligent agent by interacting with an unknown environment, and has therefore attracted more attention in economics. Furthermore, there are great application perspectives in domains such as robotics [13] or human-level intelligent game play [30].

Within ML methods, artificial intelligences are often realized by *Artificial Neural Networks* (ANNs) [6, 16, 20, 32]. The topology of ANNs is application-specific and usually chosen based on empirical evidences. Unfortunately, although two ANNs with different topologies can theoretically represent the same function [15], the most efficient topology is not readily identifiable. In particular, for sequential problems where an output depends on previous decisions, it is difficult to find such a topology.

Neuroevolution (NE) is the artificial evolution of ANNs using *Genetic Algorithms* (GAs) and explores through a search space for an ANN that performs well on a given task. It offers an alternative to statistical methods for solving control tasks that attempt to estimate the utility of actions in states of the world [2]. Studies have shown that NE is more efficient than RL methods like Q-learning

and Adaptive Heuristic Critic for several RL problems like robotic control [17, 26]. Compared to RL, a direct exploration in the search space is possible without the requirement for indirect inferences from value functions. Furthermore, problems can be dealt with effectively where gradients are difficult to compute or do not exist.

Besides conventional NE, there are *Topology and Weight Evolving Artificial Neural Network* (TWEANN) algorithms that evolve topology in addition to ANN weights. However, a downside of TWEANN algorithms is that in order to solve large-scale RL problems, the size of ANNs can grow rapidly and lead to inefficient large ANNs [25]. Thus, the evolution of ANNs has a long runtime and large memory requirements, which makes it difficult to solve specific tasks. The reason for this is that a large search effort is made within a complex search space, where there are few suitable solutions. Although efforts have been made to improve the capability of TWEANN algorithms to solve problems with large state spaces, there is still room for improvement for existing technologies [3, 23, 31]. For example, HyperNEAT uses an indirect encoding method to evolve large ANNs, but performs poorly on problems with high-level and geometrically unrelated features [3].

This paper proposes a new TWEANN algorithm called *Artificial Life Form* (ALF), which has the following main advantages over traditional TWEANN algorithms: (1) speciation via structural and semantic similarity to form better candidate solutions, (2) dynamic adaptation of the observed candidate solutions for better convergence properties, and (3) integration of solution quality into genetic reproduction to increase the probability of optimization success. The synergy of these components improves the runtime and memory usage, leading to efficient evolved ANNs. Experimental results on large-scale RL problems confirm this performance.

In the remainder of this paper, the made contributions are described in the following. The next section gives an overview of TWEANN and the resulting design challenge. Afterwards, Section 3 presents the proposed solution called ALF for these problems by describing its three components: (1) *Semantic and Structural Speciation* (SSS), (2) *Dynamic Adaptation of Population* (DAP), and (3) *Fitness-Based Genetic Operators* (FBGO). Then, the performance of ALF on large-scale RL problems is considered in Section 4, which summarizes the conducted experimental evaluations. Finally, the paper is concluded in Section 5.

2 PRELIMINARIES

This section addresses issues of evolutionary computation, ANNs and biology in an attempt to keep this work self-contained.

Algorithm 1: Possible realization of a traditional GA

Input: Fitness threshold f_t
Output: Best individual $c \in P(t)$

- 1 $t \leftarrow 0$
- 2 Initialize population $P(t)$ consisting of n individuals
- 3 **while** fitness threshold f_t is not reached **do**
- 4 Evaluate fitness $f(c) \forall c \in P(t)$
- 5 **for** $i = 1, 2, \dots, n$ **do**
- 6 Select parents c_l, c_r ($l, r \in [1, n], l \neq r$)
- 7 $c' \leftarrow \text{crossover}(c_l, c_r)$
- 8 mutate(c')
- 9 $P(t+1)_i \leftarrow c'$
- 10 **end for**
- 11 $t \leftarrow t + 1$
- 12 **end while**
- 13 **return** best $c \in P(t)$

2.1 Genetic Algorithms (GAs)

GAs are a family of computational models inspired by Darwinian natural selection. The original GA was introduced by Holland [11]. Algorithm 1 shows a simplified traditional realization of a GA, in which a population P consists of n candidate solutions (individuals), where each individual c represents a point in the search space.

Traditionally, the initial population consists of random individuals, whereby various strategy parameters like the population size are predetermined (lines 1 to 2). Each population created in generation t is called $P(t)$. The objective is to explore the search space for a solution that optimizes a performance criterion (fitness threshold) f_t (line 3). A potential solution (individual's phenotype) to a problem is encoded in a chromosome-like data structure, i. e., a string of genes, called a genome (genotype). Genes can be binary, real or correspond to any other searchable encoding scheme. The genotype is transformed into a phenotype by a specific decoding function and evaluated by a fitness function f for a task (line 4). The parents with higher fitness are then allowed to mate to reproduce. They are placed in a mating pool (line 6) and have a higher probability of reproduction. Here the hope is that the crossover of two good gene sets will result in better genes (line 7). Mutations cause random perturbations to the genes during reproduction and enable the exploration of new regions in the search space (line 8). Analogous to biological mutations, they ensure genetic diversity from one population to the next (lines 9 to 11) [29]. Finally, the best individual is returned (line 13).

In general, GAs can vary significantly. For example, the way a genotype is encoded can have a significant impact on the number of generations or result in no solution being found [26]. In addition, it is also significant how large the search space or how large the number of solutions is. Therefore, one area where ALF differs from traditional GAs is that the number of individuals or, respectively, the number of genes, can be dynamically increased and decreased at runtime, which provides more flexibility and further reduces the odds of getting stuck in local optima.

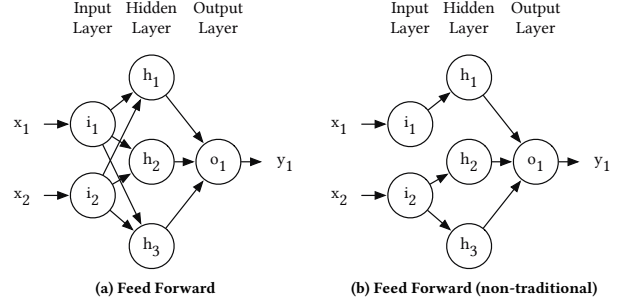


Figure 1: Fully and partially connected FFN topologies

2.2 Artificial Neural Networks (ANNs)

ANNs are computational models inspired by biological neural networks that seek to mimic the behavior and adaptive capabilities of the central nervous system [7]. They are theoretically capable of approximating any continuous function [15]. This property makes them powerful tools for control and prediction.

An ANN can be conceptualized in terms of layers and typically consists of fully connected processing units called neurons (nodes). There is an input layer with input nodes that receive data from the external world and pass their activations forward to other nodes via weighted connections. Usually, the number of input nodes corresponds to the number of features in the dataset. The last layer in an ANN is the output layer, whose outputs are used by the respective system containing the ANN. Typically, the number of output nodes corresponds to the number of classes to be predicted in the dataset. There may be hidden nodes between the input and output layers, which are located in hidden layers. Different types of networks are distinguished. Figure 1 shows ANNs, in which activation starts at the inputs and flows forwards to the outputs of the network. Such ANNs are called *Feed Forward Networks* (FFNs) [14]. Although ANNs are traditionally considered to be strictly layered and fully connected, which is shown in Figure 1a, ALF does not follow this tradition as its own ANN topologies can be evolved. Evolved ANN topologies are therefore unlikely to be fully connected, but rather partially connected as shown in Figure 1b.

Each node calculates a weighted sum of its respective inputs. The sum passes through an activation function ϕ , which squashes the activation into a certain range. For each node j , the output y_j for a passed input vector x is given by Equation (1)

$$y_j = \phi \left(\sum_i w_{ij} x_i + w_b \right), \quad (1)$$

where w_{ij} is the connection weight from node i to j and w_b is the connection weight for the bias. The activity level of the bias unit is usually 1 and the weight to a node can be positive or negative. If, e. g., there is a weak input from other nodes, the bias ensures that the unit remains active with a positive weight. This is useful if a node is frequently activated. The activation function is usually non-linear [21] and can represent, e. g., the sigmoid function $\frac{1}{1+e^{-x}}$. In this case, the activation would be squashed into the range $[0.0, 1.0]$. ANNs can also have feedback connections that flow back-




Figure 2: Fully and partially connected RNN topologies

wards instead of forwards. ANNs with feedback connections are called *Recurrent Neural Networks* (RNNs). Unlike FFNs, inputs are interdependent as part of a sequence and RNNs are therefore useful for learning temporal dependencies [5]. This type of network is shown in Figure 2, where Figure 2a represents a traditional RNN, analogous to the FFNs. ALF can also evolve recurrent structures if required, which is shown in Figure 2b. This is useful because it is not always necessary to have a recurrent connection for each hidden node. Often the temporal dependencies can be captured with only a few recurrent connections. Consequently, it would be a waste of search space to use unnecessary recurrent connections.


Traditionally, the ANN topology is determined a priori before training. The weights can then be trained using gradient descent methods like backpropagation [22], where a loss function is minimized by adjusting the weights. As a search heuristic, GAs have some advantages compared to other search methods. While gradient descent methods can get trapped in the local minima in the error surface, GAs minimize this pitfall by sampling multiple points on the error surface [26]. In particular, feedback tends to cause unexpected effects in the error surface, so that local minima are more likely to occur at suboptimal points on the surface. Additionally, training RNNs is slower and less reliable [12]. In this context, the vanishing gradient problem [10] in relation to RNNs should be mentioned. The gradient can decrease exponentially in the number of multiplications in RNNs, so that states do not contribute to the learning process. This circumstance also applies to FFNs analogously, with the exception that RNNs usually have a deeper structure. In contrast, GAs require minimal a priori knowledge about the problem domain and different parts of the search space can easily be searched simultaneously. The GAs can even search for ANNs in sparse reinforcement. In this context, the phenotype mentioned in Section 2.1 is gradually optimized by GAs.

3 ARTIFICIAL LIFE FORM (ALF)

In this work, we propose a new TWEANN algorithm called ALF, which is a solution that addresses the issue of efficiently evolving ANNs for large-scale RL problems. For this purpose, ALF integrates three main components: (1) SSS, (2) DAP, and (3) FBGO. We begin by explaining the genetic encoding used in ALF, followed by a detailed description of the components that address the issue mentioned.

	Input Layer 1			Hidden Layer 2	Output Layer 3				
φ_1	2,	$\begin{bmatrix} 1.5 & 0.2 & 0.0 \\ 0.0 & 0.0 & 0.4 \\ 0.0 & 0.0 & 1.3 \end{bmatrix}$	3,	$\begin{bmatrix} 0.0 & 1.0 & 0.0 \end{bmatrix}$	φ_2	3,	$\begin{bmatrix} 0.1 & 1.5 & 0.2 \end{bmatrix}$	φ_3	

(a) Genotype



(b) Phenotype

Figure 3: Example of genotype-phenotype mapping in ALF

3.1 Genetic Encoding

ALF uses a direct encoding scheme designed to compare and cross-over genes efficiently. Thus, a direct genotype-phenotype mapping exists for the ANN, specifying each weighted connection and layer node with an activation function that appears in the phenotype.

Genotypes are an ordered representation of genes consisting of layer connections and activation functions. A layer connection consists of a source layer id, destination layer id, and associated connection matrix. This matrix is a non-empty $(m \times n) = (a_{ij})$ matrix with n source nodes, m destination nodes, and a vector a .

Each floating-point number in the matrix at position $a_{ij} = [i \cdot n - n + j]$ ($1 \leq i \leq m$, $1 \leq j \leq n$) is the weight of the respective connection between node i and j .

In this context, we determined experimentally that this kind of implementation provides efficient cache support and optimizes memory accesses with respect to various operations such as multiplication within the ANN evaluation.

In ALF, a bias unit can be configured. The bias unit is an input that is always set to 1. If there is a connection from a bias unit to a layer, it is encoded into the connection matrices of the respective layer, whereby connections to the input layer are prohibited.

ALF has the ability to evolve topologies of different sizes. In this context, a zero entry in the connection matrix is interpreted as *no connection*. Figure 3 shows an example of a corresponding genotype-phenotype mapping. There is an input, a hidden, and an output layer as well as a bias unit b . The input layer with the id 1 contains two nodes and a layer connection to the hidden layer as well as a bias connection encoded in the first column, and a layer connection to the output layer. All possible bias connections are highlighted in italics. The hidden layer with the id 2 contains three nodes and a layer connection to the output layer. The output layer contains one node. Figure 3a shows the corresponding encoding including the activation functions of the layers. If the individual connections are considered, a zero connection can be detected between the input ($j = 2$) and the hidden node ($i = 2$), i. e., one that is not expressed in the phenotype as shown in Figure 3b. One of the

main problems of NE is the permutations problem [8], i. e., there is more than one possibility to express a solution to a weight optimization problem with an ANN. Specifically, if genotypes represent the same solution but do not have the same encoding, the cross-over is likely to reproduce offspring with defective genotypes. For example, crossing over two structurally identical but differently encoded ANNs could result in offspring in which critical information is lost or, respectively, genes are duplicated. To prevent this, there is at most one layer connection between one source and destination layer in ALF, whereby this layer connection is assigned to the source layer. The layer connections of a layer are sorted in ascending order by destination layer id. Due to the matrix data structure, at most one cell entry can exist as a connection between two nodes. Therefore, the genetic encoding in ALF is unique, allowing efficient comparisons between genes.

3.2 Semantic and Structural Speciation (SSS)

In TWEANN, it is unlikely that a new connection immediately expresses a useful function. Optimization often requires several generations, whereby the new structure does not have the necessary time to evolve. ANNs with such innovations need to be protected so that they have enough time to express these in useful ways. For this purpose, concepts of species were introduced in [11, 27].

A species is a group of individuals with similar characteristics that compete mainly within their own species. Thus, innovations are protected because they are not dominated by already optimized structures. To determine whether two individuals belong to the same species, a compatibility function δ is needed. While in [27] a topological comparison is made for their classification, we introduce the structural and semantic comparison in ALF.

Because of the solution of the permutations problem in Section 3.1, the structural similarity T can be measured efficiently. This comparison is performed by aligning the layer structures via the union of nodes. This preserves the semantics of ANNs to be compared, as no weighted connections are added. Then, a layer-wise structural comparison between ANNs can be performed by counting the shared connections E_{shared} and maximum connections E_{max} , implicitly including the nodes. Figure 4 shows an example of the structural comparison. Specifically, two ANNs N_1, N_2 , which are to be compared, are shown in Figure 4a, where N_2 is considered a smaller ANN because it has fewer layers compared to N_1 . For simplicity, the nodes are labeled accordingly. As a result of the union of nodes, in addition to the input and output layers, the hidden layer 3 of N_2 is compared with the hidden layer 3 of N_1 . Due to the unique encoding described in Section 3.1, the nodes h_2 of N_1 and N_2 are structurally identical and are merged, which can be seen in Figure 4b. Obviously, there is only one shared connection from h_2 to o_1 . The other connections are not shared. Thus, from $E_{shared} = 1$ and $E_{max} = 4$ it follows that the structural similarity is $T = \frac{E_{shared}}{E_{max}} = \frac{1}{4} = 0.25 \Rightarrow 25\%$. The semantic comparison B provides information about the similarity of the behavior of the considered ANNs. The behavior of an ANN is defined by its predictions on given inputs. The inputs can be data from a dataset or data collected at runtime. The similarity of the ANNs' predictions to be compared is determined by the correlation coefficient ρ , as shown in Equation (2)

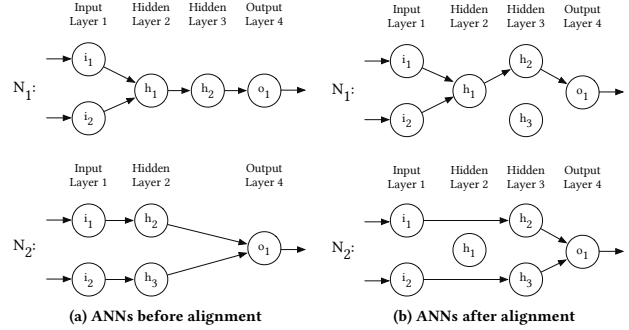


Figure 4: Example of structural comparison of ANNs

$$\rho_{X,Y} = corr(X, Y) = \frac{\sum_{i=1}^n (X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^n (X_i - \bar{X})^2} \cdot \sqrt{\sum_{i=1}^n (Y_i - \bar{Y})^2}}, \quad (2)$$

where X refers to the predictions of one ANN and Y refers to the predictions of another ANN. The correlation coefficient measures the linear dependence of the individuals. Since similarity is characterized only by correlation, anti-correlation is disregarded. For this reason, the function $B = \max(0, \rho_{X,Y})$ applies. The significance of the correlation is tested using the t-statistic [24]. Since there is not necessarily normally distributed data, we apply the central limit theorem [1] from a sample size of 30 samples. Furthermore, we require at least a 80 % confidence level for the correlation to be considered significant.

While structural species comparison protects topological diversity, semantic comparison protects diversity of behavior patterns. Both factors contribute to increasing genetic diversity. Depending on the problem, the importance of both factors can be adjusted by the coefficients c_1 and c_2 , which is shown in Equation (3)

$$\delta = \frac{T \cdot c_1 + B \cdot c_2}{c_1 + c_2}. \quad (3)$$

The compatibility measure allows a simple speciation of the population. In each generation, the offspring are classified into species. Each species is represented by the containing fittest individual (mascot) from the previous generation. An offspring is compared with each mascot and assigned to the most similar species. In this context, the compatibility threshold δ_t indicates the minimum similarity required. If $\delta \geq \delta_t$ does not apply to any species, a new species is created and the offspring becomes the mascot.

The objective of the component SSS is accordingly supporting a diversity of complex structures and behavior patterns to form better individuals in order to perform the search for a solution efficiently.

3.3 Dynamic Adaptation of Population (DAP)

In some TWEANN algorithms, the initial population consists of randomly large topologies so that topological diversity exists immediately. However, this may result in individuals with already large ANN topologies whose fitness is higher than the fitness of individuals with smaller topologies. Such ANNs initially dominate

the smaller ANNs and make the problem of finding minimal solutions more difficult, since the larger ANNs are more likely to evolve, although the smaller ANNs may also have the potential to solve the respective problem. In addition, ANNs which are topologically different but semantically identical can result. Thus, the ANN of each individual of the initial population in ALF is minimally initialized, i. e., consisting of an input and output layer with one layer connection of random weights. This reduces the probability of the dominance, but still provides sufficient initial genetic diversity.

Another problem can occur if a population size cannot change within the generations, as this may lead to stagnation of exploration in unsuitable search regions that do not advance learning. Therefore, a dynamic population is introduced in ALF, i. e., the population can increase and decrease, which is described as follows.

In ALF, fitness sharing is used for the reproduction mechanism, i. e., a species s_i contains the average fitness f_{s_i} of its assigned individuals, where $i = 1, \dots, n$ and n is the number of species. The fitness proportion of the species is denoted as $F_{s_i} = \frac{f_{s_i}}{\sum_{i=1}^n f_{s_i}}$. In addition, each species is assigned an age A_{s_i} , which identifies how many generations this very species has been in the population. The age of the oldest species is denoted as A_o . To decide whether a species is to be deleted, both factors are considered in combination, which is shown by Equation (4)

$$F_{s_i} < \frac{1}{n} \cdot c_3 \wedge A_{s_i} > A_o \cdot c_4, \quad (4)$$

where c_3 and c_4 are coefficients for adjusting the importance of these factors. Age dependency in combination with the fitness proportion penalizes older species while protecting younger species. Older species, which have a lower proportion of the population's fitness compared to other species, are treated like so-called stale species. A species is called stale if it has not improved over a predetermined number *staleness* of generations. These unsuccessful species are therefore deleted.

Furthermore, the weakest individuals of all species are deleted, whereby their number depends on F_{s_i} of the respective species and does not exceed 50 % of the population. The population is then replenished by the offspring of the remaining individuals in each species. The new population size P_{new} depends on the proportion of deleted species S_{del} , which is shown by Equation (5)

$$P_{new} = \begin{cases} \min(P_{max}, P_{old} \cdot (1 + S_{del})), & S_{del} > 0 \\ \max(P_{old} - P_{init} \cdot ebb, P_{init}), & S_{del} = 0 \end{cases}, \quad (5)$$

where *ebb* denotes an ebb factor. Thus, the greater the number of species deleted, the larger the population becomes, whereby the new population size P_{new} is limited by a predetermined value P_{max} . Therefore, more room is available for species to reproduce, so that search regions can be sampled more frequently to increase the probability of leaving insufficient search regions more rapidly. The ebb factor ensures that the population size is reduced if no species are deleted, whereby the population never becomes smaller than the initial population P_{init} . Since weaker individuals are deleted, fitter offspring resulting from reproduction still get sufficient time to solve the sub-problem. This also reduces unnecessary evaluations of weaker individuals.

The species also reproduce depending on their fitness. The number of respective offspring O_{s_i} is determined by Equation (6)

$$O_{s_i} = F_{s_i} \cdot (1 - o_{init}) \cdot (P_{new} - (P_{max} - P_{vacant})), \quad (6)$$

where P_{vacant} is the available room that can be used for the reproduction of offspring and o_{init} is the proportion of newly initialized ANNs. Within the respective species, the fittest individuals with a higher probability are then selected as parents to reproduce offspring. The remaining room is replenished with those newly initialized ANNs. Thus, even in the ongoing evolutionary process, new ANNs are provided so that possible dead ends can be overcome with a higher probability via new innovations.

For this reason, the objective of the component DAP is to improve convergence properties by overcoming insufficient search regions through increased simultaneous sampling.

3.4 Fitness-Based Genetic Operators (FBGO)

On the basis of the selected parents from the mating pool described in Section 3.3, variation operators are applied to reproduce offspring and explore the search space. These include mutation and crossover. In a standard mutation, genes are changed randomly, which can have a negative impact on the exploration in the search space. While small changes can increase the time required to explore the search space, large changes are not suitable for local optimization and can, thus, increase the time required to optimize found solutions from already explored search regions. If crossover is used, swapping genes of unequally fit individuals may result in unsuitable genes being inherited, leading to slow convergence. To counteract these problems, ALF introduces fitness-based mutation and crossover, which are described in the following.

A mutation in ALF can change both connection weights and ANN structures. A fitness-based mutation rate is calculated for all mutations, which determines how often ALF attempts to perform a respective mutation. The mutation rate m_r is defined as Equation (7)

$$m_r = \min \left(\max \left(1, \left(\left((1 + m_o) - \frac{F}{f_t} \right) \cdot m_a \right) \right), m_a \right), \quad (7)$$

where m_o is a predefined mutation offset. On the one hand, in relation to the fitness threshold f_t , m_o indicates up to which proportion of the current fitness F of an individual the maximum number of attempts m_a is performed. On the other hand, it ensures that m_r cannot become 0, so that mutations still take place even with high fitness values. A connection weight a_{ij} is changed by adding a random number from the normal distribution $\mathcal{N} \left(0, 1.25 - \frac{F}{f_t} \right)$. The fitter an individual is, the smaller changes are made, i. e., when F is low, compared to f_t , the search process has more influence and when F is high, the optimization process is strengthened. Based on the genetic encoding (see Section 3.1), the *weight mutation* creates, changes or deletes a connection. In order to change connections significantly more often than to add layers or nodes, it is determined with regard to the *node* and *layer mutation* in addition to m_r , whether such a respective mutation should be performed. The corresponding probability calculation P_{mut} is defined in Equation (8)

$$P_{mut} = \min \left(1 + m_o - \frac{\frac{E}{E_{full}} + \frac{F}{f_t}}{2}, 1 \right), \quad (8)$$

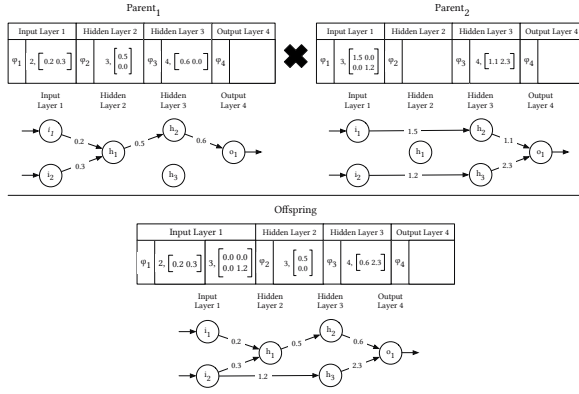


Figure 5: Example of crossing over genes of parents

where E is the current number of connections and E_{full} is the number of possible connections (assuming the ANN is fully connected). The proportion of connections to possible connections describes the potential for new connections, i. e., new connection possibilities (nodes or layers) become more likely to be added if $\frac{E}{E_{full}}$ is closer to 1. *Node mutation* randomly adds nodes into hidden layers, determining the number of nodes V_{mut} as shown in Equation (9)

$$V_{mut} = \left\lceil V_{max} - \frac{F}{f_t} \cdot V_{max} \right\rceil, \quad (9)$$

where V_{max} is predetermined and is the maximum number of nodes to be added. The *layer mutation* inserts a new layer connection or a new hidden layer. For this purpose, a layer is randomly selected and checked whether a layer connection can still be added. If not, a new hidden layer is appended. The number of nodes of the new layer is calculated as in Equation (9).

Mutations will gradually increase the size of genotypes in ALF, resulting in genotypes of different sizes. In order to be able to compare such layer structures, ALF uses layer alignment by node union as described in Section 3.2. Based on the solution of the permutations problem as described in Section 3.1, it is, thus, possible to identify gene matches efficiently. In ALF, there is a hyper-parameter P_{cross} that indicates the probability of a crossover between selected individuals. When crossing over, fitness has an impact on the inheritance. Therefore, before crossing over, the respective fitness proportion is determined in relation to the fitness sum of both parents. The fitness proportion is interpreted as a probability during the gene comparison. Accordingly, when matching genes are compared, the gene of the fitter parent is more likely to be inherited. When disjoint genes are considered, a gene of the fitter parent is more likely to be inherited, while a gene of the weaker parent is more likely to be not inherited. Consider Figure 5, which shows an example of the crossover used in ALF assuming that parent $Parent_1$ is fitter than $Parent_2$. While all disjoint genes are inherited from $Parent_1$, the gene from node i_1 to h_2 is not inherited from $Parent_2$. For the matching gene from h_2 to o_1 , the higher fitness of $Parent_1$ causes its gene to be inherited. By fitness-based adding of new genes to the population and crossing over genotypes, the objective of the component FBGO is to increase the probability of optimization success of individuals.

4 EXPERIMENTAL RESULTS

In this section, ALF is analyzed empirically. For this purpose, we implemented the components described in the previous section in C++17 and measured their performances in terms of the average time in minutes (t in min), average number of generations (G), connections (E), and hidden nodes (H) required by applying them to large-scale RL problems. The analysis of ALF aims at answering the following questions: (1) Does the symbiosis of the proposed components of ALF lead to an effective evolution of necessary ANNs? (2) Does ALF find solutions efficiently? The first question establishes that the symbiosis of ALF’s components indeed reduces the memory usage and runtime. To answer the second question, in this paper we choose the TWEANN algorithm *NeuroEvolution of Augmenting Topologies* (NEAT) [27] as test algorithm for the evaluation, because it has already proven to be successful in complex control learning tasks [26] and RAM-based Atari games, and was more reliable than other TWEANN algorithms like HyperNEAT [9]. Compared to other NE algorithms, NEAT is characterised in particular by topology evolution and the protection of innovations via speciation by topological comparisons of genotypes. We limit ourselves to NE because the focus is on achieving better performance of evolving ANNs and NE algorithms demonstrated similar performance in games compared to RL methods such as deep Q-learning in the literature, whereby RL methods required significantly more training time [28].

4.1 SMB as a Benchmark Task

Super Mario Bros. (SMB) is a game for the *Nintendo Entertainment System* (NES) that is divided into several game worlds with different levels (problems). In each level, the protagonist (Mario) starts from a specific location and his objective is to reach the flagpole or axe in a certain time, which leads to the next level. Occasionally, a level contains hidden warp pipes, which let him skip some levels. In addition, there are various challenges such as canyons, enemies or dead ends that Mario has to overcome. Another game element is the loop command object, an invisible tile that resets Mario’s position by several screens if he does not follow a right path. If Mario dies by walking into an enemy or falling into a canyon, he loses one life and has to start from the beginning of the level. If he reaches a checkpoint, it becomes the new starting point of the level.

The criterion for success on this task was the completion of the levels considered. The task was suitable to test ALF’s abilities because of several reasons. SMB is a sequential and unlabeled problem, which is not linearly separable. Thus, hidden nodes have to be evolved. Furthermore, the task of reaching the flagpole or axe is PSPACE-complete [4], whereby a large solution space with few suitable solutions exists unlike in other RL problems like cart-pole [19].

To consider many game elements, three game worlds were studied and considered as large-scale RL problems: 4-2, 4-4, and 7-2. World 4-2 contains canyons, enemies and passages, in which Mario has to wait a certain time. Moreover, this level contains a possible dead end, where Mario has to go back and jump at the correct time to overcome a high obstacle. World 4-4 contains the mentioned loop command objects. For both levels, a fitness function was defined as in Equation (10)

$$f(d, q, x) = d + q + x, \quad (10)$$

where d is the game progress, q is the scroll amount, and x denotes the x -coordinate of Mario’s position. While the fitness threshold $f_t = 3,163$ of world 4-2 referred to the x -coordinate of the flagpole, $f_t = 2,770$ of world 4-4 corresponded to the x -coordinate of the axe.

World 7-2 is a water level, where Mario can also move vertically. Therefore, the y -coordinate of Mario’s position was additionally set in relation to his x -coordinate, which was defined by Equation (11)

$$f(d, q, x, y) = d + q + 255x + (255 - y). \quad (11)$$

The fitness threshold of world 7-2 corresponded to the x -coordinate of the flag pole, i. e., $f_t = 3,404$.

Regarding the activation function, we oriented ourselves on the sigmoid variant of NEAT [27] and modified it to support inhibitory signals. The function is shown in Equation (12). It squashes the incoming values v into the range $[-1.0, 1.0]$. To reduce the dimension of the learning problem, a bias unit was configured.

$$\phi(v) = \text{sig}_m(v) = \frac{2}{1 + e^{-4.9v}} - 1 \quad (12)$$

The NES has a frame consisting of 184,320 pixels. To handle this large input space, we compressed it into 80 inputs, which we considered as sufficiently small for our empirical study. We read out a field of view around Mario’s position from the memory, whereby each enemy and game element was encoded differently to distinguish them. A frame is divided into blocks and tiles. A tile consists of 8×8 pixels and a block consists of 2×2 tiles. We chose a field of view of 10×8 blocks, i. e., we chose an area of $x = [-2, 8]$ and $y = [-5, 3]$ blocks relative to Mario’s position. This number of blocks was determined through experiments and observations of the level structures of SMB. Each block represented a feature, which was determined by the mean of the tile values within the block. Enemies were encoded as negative values, while other game elements were encoded as positive values. Mario was encoded as 0. Each encoded value was scaled to a range of $[-1.0, 1.0]$ before being passed to the ANN in order to reduce feature dominance and, therefore, possible generalization errors. Moreover, the number of output nodes was set to 6, whereby each node represented a button on the NES controller. These were the buttons that the ANN had to press or release after each processed frame.

All experiments were conducted on an Intel Xeon E3-1230 v6 with 3.2 GHz and 16 GB of main memory running Fedora 22. Because of the complexity of SMB, 10 runs were performed and the respective average, possible extremities, and divergences were determined. The timeout to finish a level was set to 2 days.

4.2 Verification of ALF’s Components

In this section, it is verified whether the introduced components of ALF are reliable and their symbiosis leads to reduced runtime and memory usage. Therefore, ALF was applied to the SMB world 4-2.

4.2.1 Hyper-parameter Settings. For all experiments, the same experimental settings were used. However, the parameters were not adapted to SMB. Table 1 shows the most important hyper-parameters used in ALF, which were determined experimentally and are described in the following. The initial population size con-

Table 1: Hyper-parameter settings for ALF

Hyper-parameter	ALF
Initial population size (P_{init})	150
Maximum population size (P_{max})	300
Ebb factor (ebb)	0.05
Mutation offset (m_o)	0.1
Maximum mutation attempts (m_a)	5
Maximum added nodes (V_{max})	5
Crossover probability (P_{cross})	0.7
Staleness ($staleness$)	15
Oldest species age coefficient (c_4)	0.5
Species fitness proportion coefficient (c_3)	0.4
Newly initialised ANNs proportion (o_{init})	0.1
Semantic coefficient (c_2)	0.75
Structural coefficient (c_1)	0.25
Compatibility threshold (δ_t)	0.3

sists of 150 individuals, with a maximum size of 300 individuals. This provides sufficient room for more species, allowing for greater genetic diversity. If no species is deleted, the population shrinks by 5 %, whereby the population size never gets smaller than the initial population size. Otherwise, the population grows in relation to the number of deleted species to provide additional individuals. This ensures that the exploration of the search space is never reduced if there is an approximation to an insufficient local optimum.

Until a fitness proportion of 10 % in relation to the fitness threshold is reached, the maximum number of mutation attempts is performed on the genes. Otherwise, the number of attempts is decreased in relation to the fitness and number of connections, in order to approximate to a local optimum more purposefully. Since connections are to be added more frequently than nodes, a proportionally low maximum number of nodes is chosen with respect to the mutation. The probability of reproducing an offspring without any crossover is 30 %, which is common for GAs.

If the maximum fitness of a species does not improve within 15 generations, the species is considered as stale and is deleted. The species is also deleted if it is older than half of the oldest species and its fitness proportion is smaller than 40 %, assuming the fitness is equally distributed. Thus, a species has enough time to evolve, but is penalized if it does not. One exception is a species that contains the best individual to prevent degradation of the best fitness in the population. The weakest individuals of each species are deleted, whereby the number of deleted individuals depends on the adapted species’ fitness. However, the number of deleted individuals never exceeds 50 % of the population size and the best individual of a species is never deleted. Each new population is replenished with an offspring proportion of 90 %. The behavior of the ANNs is considered more important than their structural information. Therefore, the semantic coefficient is set to 75 %, while the structural coefficient is set to 25 %. In this context, it is required that an offspring, compared to a mascot, has to achieve a minimum similarity of 30 % to be assigned. Otherwise, a new species is created.

4.2.2 Results. Based on the initial run (*no component*) without using any of the addressed components (1) SSS, (2) DAP, and

Table 2: Comparison between component symbioses of ALF in terms of runtime and ANN sizes to solve SMB world 4-2

Symbiosis of ALF's components	SMB			
	World 4-2			
	t in min	G	E	H
No component	1,801	603	4,288	1,316
SSS	1,363	334	3,522	1,021
DAP	1,789	572	4,109	1,193
FBGO	1,456	478	3,768	687
SSS+DAP	1,073	241	3,004	980
SSS+FBGO	1,106	256	2,703	389
DAP+FBGO	1,194	328	3,235	546
SSS+DAP+FBGO	978	194	2,445	304

t in min	Time in minutes
G	Number of generations needed
E	Number of evolved connections of the ANNs
H	Number of evolved hidden nodes of the ANNs
SSS	Semantic and Structural Speciation
DAP	Dynamic Adaptation of Population
FBGO	Fitness-Based Genetic Operators

(3) FBGO, we aimed at determining whether the SMB world 4-2 can be solved faster with fewer average connections and hidden nodes in fewer generations. The results are summarized in Table 2 and described in the following.¹ ALF is consistent in finding a solution. In particular, the symbiosis of the components results in less time, generations, connections, and hidden nodes required on average to solve this problem.

On average, the worst case without using any component needs 1,801 minutes, 603 generations, 4,288 connections, and 1,316 hidden nodes, while the best symbiosis (SSS+DAP+FBGO) requires 978 minutes, 194 generations, 2,445 connections, and 304 hidden nodes. This reduces the ANN size by about 50.95%. Furthermore, the problem is solved almost 2 times faster. In particular, the addition of FBGO significantly reduces the number of hidden nodes or, respectively the proportion between connections and hidden nodes. Therefore, e. g., the number of hidden nodes in SSS+FBGO is reduced by about 31.94% and the number of hidden nodes in DAP+FBGO is reduced by about 28.69%.

In the context of SSS, we observed that a species, which went through the bonus area, was evolved. Individuals of this species went through a bonus pipe and were able to solve the problem faster. With respect to DAP, we observed that population size increased when the population had to solve complex sub-problems like crossing canyons. Due to more individuals, these sub-problems could be solved faster, especially in symbioses with FBGO or SSS.

In conclusion, ALF solves SMB world 4-2 without any difficulties. Specifically, by adding the mentioned components (in symbioses), this problem is solved faster with less memory usage.

4.3 Evaluation of ALF's Performance

To obtain meaningful performance comparisons with NEAT and to determine whether ALF is more efficient in finding solutions, SMB worlds 4-2, 4-4 and 7-4 are considered. First, the experimental setup is described, followed by the analysis of ALF's performance.

¹Note when SSS was deactivated, each individual was assigned to a separate species.

Table 3: Comparison between NEAT and ALF in terms of runtime and ANN sizes to solve different problems

SMB	Algorithm							
	NEAT				ALF			
	t in min	G	E	H	t in min	G	E	H
World 4-2	1,438	387	3,873	1,159	978	194	2,445	304
World 4-4	1,373	372	3,621	1,106	889	157	2,219	339
World 7-2	1,894	486	4,305	1,438	1,454	273	3,252	483

t in min	Time in minutes
G	Number of generations needed
E	Number of evolved connections of the ANNs
H	Number of evolved hidden nodes of the ANNs

4.3.1 Hyper-parameter Settings. For the following experiments we used the same hyper-parameter settings for ALF as described in Section 4.2.1. Moreover, we set up NEAT as described in [9] and adjusted the hyper-parameters, i. e., the same available genetic parameters, such as population size, were matched accordingly to allow a representative comparison.

4.3.2 Results. Next, we compare ALF with NEAT in order to determine whether ALF can find solutions efficiently. The results are summarized in Table 3 and clearly confirm that the components proposed satisfy the objectives of this work. First, the proposed algorithm is capable of effectively solving the large-scale RL problems mentioned. Second, it addresses one of the known drawbacks of existing solutions: the evolution of inefficient ANNs for large-scale RL problems.

In this context, ALF outperforms NEAT in terms of runtime, generations, connections and hidden nodes required on each game world. In total, ALF solves the problems about 461 minutes faster on average and requires 207 generations less compared to NEAT. In addition, the ANN sizes are reduced by about 41.67% on average.

Overall, the proposed algorithm is able to solve large-scale RL problems effectively, leading to efficient ANNs.

5 CONCLUSION

This work focused on one of the main challenges in NE – the efficient solving of large-scale RL problems. The TWEANN algorithm ALF presented here can handle large-scale RL problems via the synergy of different components – structural and semantic speciation, dynamic adaptation of observed candidate solutions, and integration of solution quality into genetic reproduction. Experimental results for large-scale RL problems confirmed that the proposed components are able to evolve efficient ANNs effectively. Future work will include optimizing the components described and extending the performance comparisons with other test algorithms and benchmark tasks.

REFERENCES

- [1] Carolyn Anderson. 2010. *Central Limit Theorem*. <https://doi.org/10.1002/9780470479216.corpsy0160>
- [2] Kai Arulkumaran, Marc Deisenroth, Miles Brundage, and Anil Bharath. 2017. A Brief Survey of Deep Reinforcement Learning. *IEEE Signal Processing Magazine* 34 (08 2017). <https://doi.org/10.1109/MSP.2017.2743240>
- [3] Thomas Berg and Shimon Whiteson. 2013. Critical factors in the performance of HyperNEAT. *GECCO 2013 - Proceedings of the 2013 Genetic and Evolutionary Computation Conference*, 759–766. <https://doi.org/10.1145/2463372.2463460>

- [4] Erik Demaine, Giovanni Viglietta, and Aaron Williams. 2016. Super Mario Bros. Is Harder/Easier than We Thought. (06 2016).
- [5] K.-L. Du and M.N.s Swamy. 2014. *Recurrent Neural Networks*. 337–353. https://doi.org/10.1007/978-1-4471-5571-3_11
- [6] A. Graves and N. Jaitly. 2014. Towards end-to-end speech recognition with recurrent neural networks. *31st International Conference on Machine Learning, ICML 2014* 5 (01 2014), 1764–1772.
- [7] Enzo Grossi and Massimo Buscema. 2008. Introduction to artificial neural networks. *European journal of gastroenterology & hepatology* 19 (01 2008), 1046–54. <https://doi.org/10.1097/MEG.0b013e3282f198a0>
- [8] Peter Hancock. 1997. Genetic Algorithms and permutation problems: a comparison of recombination operators for neural net structure specification. (09 1997).
- [9] Matthew Hausknecht, Joel Lehman, Risto Miikkulainen, and Peter Stone. 2014. A Neuroevolution Approach to General Atari Game Playing. *Computational Intelligence and AI in Games, IEEE Transactions on* 6 (12 2014), 355–366. <https://doi.org/10.1109/TCAIG.2013.2294713>
- [10] Sepp Hochreiter. 1998. The Vanishing Gradient Problem During Learning Recurrent Neural Nets and Problem Solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 6 (04 1998), 107–116. <https://doi.org/10.1142/S0218488598000094>
- [11] John H. Holland. 1992. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence*. MIT Press, Cambridge, MA, USA.
- [12] Fakultit Informatik, Y. Bengio, Paolo Frasconi, and Jfirgen Schmidhuber. 2003. Gradient Flow in Recurrent Nets: the Difficulty of Learning Long-Term Dependencies. *A Field Guide to Dynamical Recurrent Neural Networks* (03 2003).
- [13] Jens Kober, J. Bagnell, and Jan Peters. 2013. Reinforcement Learning in Robotics: A Survey. *The International Journal of Robotics Research* 32 (09 2013), 1238–1274. <https://doi.org/10.1177/0278364913495721>
- [14] B.K. Lavine and T.R. Blank. 2010. Feed-Forward Neural Networks. *Comprehensive Chemometrics* 3 (01 2010), 571–586. <https://doi.org/10.1016/B978-044452701-1.00026-0>
- [15] Grzegorz Lewicki and Giuseppe Marino. 2004. Approximation by Superpositions of a Sigmoidal Function. *Appl. Math. Lett.* 17 (12 2004), 1147–1152. <https://doi.org/10.1016/j.aml.2003.11.006>
- [16] Shujie Liu, Nan Yang, Mu Li, and Ming Zhou. 2014. A Recursive Recurrent Neural Network for Statistical Machine Translation. *52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014 - Proceedings of the Conference* 1, 1491–1500. <https://doi.org/10.3115/v1/P14-1140>
- [17] D. E. Moriarty and R. Miikkulainen. 1997. Forming Neural Networks Through Efficient and Adaptive Coevolution. *Evolutionary Computation* 5, 4 (1997), 373–399. <https://doi.org/10.1162/evco.1997.5.4.373>
- [18] Amir Mosavi, Yaser Faghan, Pedram Ghamisi, Puhong Duan, Sina Ardabili, Ely Salwana, and Shahab Band. 2020. Comprehensive Review of Deep Reinforcement Learning Methods and Applications in Economics. 8 (09 2020), 1640. <https://doi.org/10.3390/math8101640>
- [19] Savinay Nagendra, Nikhil Podila, Rashmi Ugarakhod, and Koshy George. 2017. Comparison of Reinforcement Learning algorithms applied to the Cart Pole problem. 26–32. <https://doi.org/10.1109/ICACCL2017.8125811>
- [20] Cong Nguyen, Hoang Dinh Thai, Shimin Gong, Dusit Niyato, Ping Wang, Ying-Chang Liang, and Dong In Kim. 2019. Applications of Deep Reinforcement Learning in Communications and Networking: A Survey. *IEEE Communications Surveys & Tutorials* PP (05 2019), 1–1. <https://doi.org/10.1109/COMST.2019.2916583>
- [21] A Olgac and Bekir Karlik. 2011. Performance Analysis of Various Activation Functions in Generalized MLP Architectures of Neural Networks. *International Journal of Artificial Intelligence And Expert Systems* 1 (02 2011), 111–122.
- [22] Mohammed Otair and A. Salameh Walid. 2006. Efficient Training of Backpropagation Neural Networks. *Neural Network World* 16 (01 2006).
- [23] Yiming Peng, Gang Chen, Harman Singh, and Mengjie Zhang. 2018. NEAT for large-scale reinforcement learning through evolutionary feature learning and policy gradient search. 490–497. <https://doi.org/10.1145/3205455.3205536>
- [24] Patrick Schober, Christa Boer, and Lothar Schwarte. 2018. Correlation Coefficients: Appropriate Use and Interpretation. *Anesthesia & Analgesia* 126 (02 2018), 1. <https://doi.org/10.1213/ANE.0000000000002864>
- [25] Kenneth Stanley, David D’Ambrosio, and Jason Gauci. 2009. A Hypercube-Based Encoding for Evolving Large-Scale Neural Networks. *Artificial life* 15 (02 2009), 185–212. <https://doi.org/10.1162/artl.2009.15.2.15202>
- [26] Kenneth Stanley and Risto Miikkulainen. 2003. Competitive Coevolution through Evolutionary Complexification. *Journal of Artificial Intelligence Research* 21 (02 2003). <https://doi.org/10.1613/jair.1338>
- [27] K. O. Stanley and R. Miikkulainen. 2002. Evolving Neural Networks through Augmenting Topologies. *Evolutionary Computation* 10, 2 (2002), 99–127. <https://doi.org/10.1162/106365602320169811>
- [28] Felipe Such, Vashisht Madhavan, Edoardo Conti, Joel Lehman, Kenneth Stanley, and Jeff Clune. 2017. Deep Neuroevolution: Genetic Algorithms Are a Competitive Alternative for Training Deep Neural Networks for Reinforcement Learning. (12 2017).
- [29] Scott Thede. 2004. An introduction to genetic algorithms. *Journal of Computing Sciences in Colleges* 20 (10 2004).
- [30] Ruben Torrado, Philip Bontrager, Julian Togelius, Jialin Liu, and Diego Perez-Liebana. 2018. Deep Reinforcement Learning for General Video Game AI. 1–8. <https://doi.org/10.1109/CIG.2018.8490422>
- [31] Shimon Whiteson, Peter Stone, Kenneth Stanley, Risto Miikkulainen, and Nate Kohl. 2005. Automatic feature selection in neuroevolution. 1225–1232. <https://doi.org/10.1145/1068009.1068210>
- [32] Libo Zhang, Tiejian Luo, Fei Zhanga, and Yanjun Wu. 2018. A Recommendation Model Based on Deep Neural Network. *IEEE Access* PP (01 2018), 1–1. <https://doi.org/10.1109/ACCESS.2018.2789866>