Evolving programs to build artificial neural networks

Julian F. Miller, Dennis G. Wilson, and Sylvain Cussat-Blanc

Abstract In general, the topology of Artificial Neural Networks (ANNs) is human-engineered and learning is merely the process of weight adjustment. However, it is well known that this can lead to sub-optimal solutions. Topology and Weight Evolving Artificial Neural Networks (TWEANNs) can lead to better topologies however, once obtained they remain fixed and cannot adapt to new problems. In this chapter, rather than evolving a fixed structure artificial neural network as in neuroevolution, we evolve a pair of programs that *build* the network. One program runs inside neurons and allows them to move, change, die or replicate. The other is executed inside dendrites and allows them to change length and weight, be removed, or replicate. The programs are represented and evolved using Cartesian Genetic Programming. From the developed networks multiple traditional ANNs can be extracted, each of which solves a different problem. The proposed approach has been evaluated on multiple classification problems.

1 Introduction

Artificial neural networks (ANNs) were first proposed seventy-five years ago [45]. Yet, they remain poor caricatures of biological neurons. ANNs are almost always

---

Julian F. Miller
University of York, Heslington, York, YO10 5DD, UK e-mail: julian.miller@york.ac.uk
It is with great pleasure that I offer this article in honour of Susan Stepney's 60th birthday. Susan has been a very stimulating and close colleague for many years.

Dennis G. Wilson
University of Toulouse, IRIT - CNRS - UMR5505, 21 allee de Brienne, Toulouse, France 31015
e-mail: dennis.wilson@irit.fr

Sylvain Cussat-Blanc
University of Toulouse, IRIT - CNRS - UMR5505, 21 allee de Brienne, Toulouse, France 31015
e-mail: sylvain.cussat-blanc@irit.fr

static arrangements of artificial neurons with simple activation functions. Learning is largely considered to be the process of adjusting weights to make the behaviour of the network conform as closely as possible to a desired behaviour. We refer to this as the "synaptic dogma". Indeed, there is abundant evidence that biological brains learn through many mechanisms [32] and in particular, changes in synaptic strengths are at best a minor factor in learning since synapses are constantly pruned away and replaced by new synapses [66]. In addition, restricting learning to weight adjustment leads immediately to the problem of so-called catastrophic forgetting. This is where retraining an ANN on a new problem causes the previously learned behaviour to be disrupted if not forgotten [19, 44, 55]. One approach to overcoming catastrophic forgetting is to create developmental ANNs which, in response to environmental stimulus (i.e. being trained on a new problems), grow a new sub-network which integrates with the existing network. Such new networks could even share some neurons with pre-existing networks.

In contrast to the synaptic dogma in ANNs, in neuroscience it is well-known that learning is strongly related to structural changes in neurons. Examples of this are commonplace. Mice reared in the dark and then placed in the light develop new dendrites in the visual cortex within days [78]. Animals reared in complex environments where active learning is taking place have an increased density of dendrites and synapses [37, 38]. Songbirds in the breeding season increase the number, size and spacing of neurons, where the latter is caused by increases in dendritic arborization [74]. It is also well-known that the hippocampi of London taxi drivers are significantly larger relative to those of control subjects [43]. Rose even argues that after a few hours of learning the brain is permanently altered [60]. Another aspect supporting the view that structural changes in the brain are strongly associated with learning, is simply that the most significant period of learning in animals happens in infancy, when the brain is developing [8].

There have been various attempts to create developmental ANNs and we review past work in Sect 2. Although many interesting past approaches have been presented, generally the work has not continued. When one considers the enormous potential of developmental neural networks there needs to be sustained and long-term research on a diversity of approaches. In addition there is a need for such approaches to be applied in a variety of real-world applications.

2 Related work on the development of ANNs

Although non-developmental in nature, a number of methods have been devised which under supervision gradually augment ANNs by adding additional neurons or join trained ANNs together via extra connections. 'Constructive neural networks' are traditional ANNs which start with a small network and add neurons incrementally while training error is reduced [13, 18]. Modular ANNs use multiple ANNs each of which has been trained on a sub-problem and these are combined by a human expert [64]. Both of these approaches could be seen as a form of human engi-

neered development. More recent approaches adjust weighted connections between trained networks on sub-problems [62, 72]. Aljundi et al. have a set of trained ANNs for each task (experts) and use an additional ANN as a recommender as to which expert to use for a particular data instance [1].

A number of authors have investigated ways of incorporating development to help construct ANNs [41, 70]. Researchers have investigated a variety of genotype representations at different levels of abstraction.

Cangelosi et al. defined genotypes which were a mixture of variables, parameters, and rules (e.g. cell type, axon length and cell division instructions) [6]. The task was to control a simple artificial organism. Rust et al constructed a binary genotype which encoded developmental parameters that controlled the times at which dendrites could branch and how the growing tips would interact with patterns of attractants placed in an environment [61]. Balaam investigated controlling simulated agents using a two-dimensional area with chemical gradients in which neurons were either sensors, affectors, or processing neurons according to location [3]. The neurons were defined as standard continuous time recurrent neural networks (CTRNNS). The genotype was effectively divided into seven chromosomes each of which read the concentrations of the two chemicals and the cell potential. Each chromosome provided respectively the neuron bias, time constant, energy, growth increment, growth direction, distance to grow and new connection weight.

A variety of grammar-based developmental methods for building ANNs have been proposed in the literature. Kitano evolved matrix re-writing rules to develop an adjacency matrix defining the structure of a neural network [36]. He used backpropagation to adjust the connection weights. He applied the technique to encoder-decoder problems of various sizes. Kitano claimed that his method produced superior results to direct methods. However, it must be said, it was later shown in a more careful study by Siddiqi and Lucas, that the two approaches were of equal quality [65]. Belew [4] evolved a two-dimensional context sensitive grammar that constructed polynomial networks (rather than ANNs) for time-series prediction. Gradient descent was used to refine weights. Genotypes were variable length with a variety of mutation operators. He found a gene-doubling mutation to be critically important. Gruau devised a grammar-based approach called cellular encoding in which ANNs were developed using graph grammars [22, 23]. He evaluated this approach on hexapod robot locomotion and pole-balancing. Kodjabachian and Meyer used a "geometry-orientated" variant of cellular encoding to develop recurrent neural networks to control the behaviour of simulated insects [39]. Drchal and Šnorek [10] replaced the tree-based cellular encoding of Gruau with alternative genetic programming methods, namely grammatical evolution (GE) [63] and gene expression programming (GEP) [16]. They also investigated the use of GE and GEP to evolve edge encoded [42] development trees. They evaluated the developed ANNs on the two-input XOR problem. Jung used a context-free grammar to interpret an evolved gene sequence [31]. When decoded it generates 2D spatially modular neural networks. The approach was evaluated on predator-prey agents using a coevolution. Floreano and Urzelai [17] criticised developmental methods that developed fixed weights. They used a matrix re-writing method inspired by Kitano, but each connection ei-

ther had a fixed weight or the neuron could choose one of four synaptic adjustment rules (i.e. plasticity). After development stopped synaptic adjustment rules were applied to all neuron connections. They applied their method to robot controllers and found that synaptic plasticity produced ANNs with better performance than fixed weight networks. In addition they were more robust in changed environments.

Jacobi presented a low-level approach in which cells used artificial genetic regulatory networks (GRNs). The GRN produced and consumed simulated proteins that defined various cell actions (protein diffusion movement, differentiation, division, threshold). After a cellular network had developed it was interpreted as a neural network [30]. Eggenberger also used an evolved GRN [12]. A neural network phenotype was obtained by comparing simulated chemicals in pairs of neurons to determine if the neurons are connected and whether the connection is excitatory or inhibitory. Weights of connections were initially randomly assigned and Hebbian learning used to adjust them subsequently. Astor and Adami devised a developmental ANN model known as Norgev (Neuronal Organism Evolution) which encoded a form of GRN together with an artificial chemistry (AC), in which cells were predefined to exist in a hexagonal grid. Genes encoded conditions involving concentrations of simulated chemicals which determine the level of activation of cellular actions (e.g. grow axon or dendrite, increase or decrease weight, produce chemical) [2]. They evaluated the approach on a simple artificial organism. In later work, Hampton and Adami showed that neurons could be removed and the developmental programs would grow new neurons and recover the original functionality of the network (it computed the NAND function) despite having a non-deterministic developmental process [24]. Yerushalmi and Teicher [80] presented an evolutionary cellular development model of spiking neurons. Inside the cells was a bio-plausible genetic regulatory network which controls neurons and their dendrites and axons in 2D space. In two separate experiments the GRN in cells was evolved to produce: (a) specific synaptic plasticity types (Hebbian, AntiHebbian, Non-Hebbian and Spike-Time Dependent Plasticity) and (b) simple organisms that had to mate. In the latter, they examined the types of synaptic plasticity that arose.

Federici used a simple recursive neural network as a developmental cellular program [14]. In his model, cells could change type, replicate, release chemicals or die. The type and metabolic concentrations of simulated chemicals in a cell were used to specify the internal dynamics and synaptic properties of its corresponding neuron. The weight of a connection between cells was determined by the difference in external chemical concentrations produced by the two cells. The position of the cell within the organism is used to produce the topological properties of a neuron: its connections to inputs, outputs and other neurons. From the cellular phenotype, Federici interpreted a network of spiking neurons to control a Khepera robot.

Roggen et al. devised a highly simplified model of development that was targeted at electronic hardware [59]. Circuits were developed in two phases. Diffusers are placed in a cellular grid and diffuse signals (analogous to chemicals) to local neighbours. There can be multiple signal types and they do not interact with each other. At the same time as the diffusion phase an expression phase determines the function of the cells by matching signal intensities with a signal-function expression

table. A genetic algorithms is used to evolve the positions of diffusing cells (which are given a maximum signal) and the expression table. By interpreting the functionalities as various kinds of spiking neurons with particular dendritic branches, the authors were able develop spiking neural networks. Connection weights were fixed constants with a sign depending on whether the pre-synaptic neuron is excitatory or inhibitory. They evaluated the approach on character recognition and robot control.

Some researchers have studied the potential of Lindenmeyer systems for developing artificial neural networks. Boers and Kuiper adapted L-systems to develop artificial feed-forward neural networks [5]. They found that this method produced more modular neural networks that performed better than networks with a predefined structure. They showed that their method could produce ANNs for solving problems such as the XOR function. Hornby and Pollack evolved L-systems to construct complex robot morphologies and neural controllers [27, 26].

Downing developed a higher-level, neuroscience-informed approach which avoided having to handle axonal and dendritic growth, and maintained important aspects of cell signaling, competition and cooperation of neural topologies [9]. He adopted ideas from Deacon's Displacement Theory [7] which is built on Edelman's Darwinistic view of neurogenesis, known as "The Theory of Neural Group Selection" [11]. In the latter, neurons will only reach maturity if they grow axons to, and receive axons from, other neurons. Downing's method has three phases. The first (translation) decodes the genotype which defines one or more neuron groups. In the second phase (displacement) the sizes and connectivity of neuron groups undergo modification. In the final stage (instantiation) populations of neurons and their connections are established. He applied this technique to the control of a multi-limbed starfish-like animat.

Khan and Miller created a complex developmental neural network model that evolved seven programs each representing various aspects of idealised biological neurons [33]. The programs were represented and evolved using Cartesian Genetic Programming (CGP) [50]. The programs were divided into two categories. Three of the encoded chromosomes were responsible for 'electrical' processing of the 'potentials'. These were the dendrite, soma and axo-synapse chromosomes. One chromosome was devoted to updating the weights of dendrites and axo-synapses. The remaining three chromosomes were responsible for updating the neural variables for the soma (health and weight), dendrites (health, weight and length) and axo-synapse (health, length). The evolved developmental programs were responsible for the removal or replication of neural components. The model was used in various applications: intelligent agent behaviour (wumpus world), checkers playing, and maze navigation [34, 35].

Although not strictly developmental, Koutnik et al. [40] investigated evolving a compression of the ANN weight matrix by mapping it to a real-valued vector of Fourier coefficients in the frequency domain. This idea reduces the dimensionality of the weight space by ignoring high-frequency coefficients, as in lossy image compression. They evaluated the merits of the approach on ANNs solving pole-balancing, ball throwing and octopus arm control. They showed that approach found solutions in significantly fewer evaluations than evolving weights directly.

Stanley introduced the idea of using evolutionary algorithms to build neural networks constructively (called NEAT). The network is initialised as a simple structure, with no hidden neurons consisting of a feed-forward network of input and output neurons. An evolutionary algorithm controls the gradual complexification of the network by adding a neuron along an existing connection, or by adding a new connection between previously unconnected neurons [67]. However, using random processes to produce more complex networks is potentially very slow. It also lacks biological plausibility since natural evolution does not operate on aspects of the brain directly. Later Stanley introduced an interesting and popular extension to the NEAT approach called HyperNEAT [69] which uses an evolved generative encoding called a Compositional Pattern Producing Network (CPPN) [68]. The CPPN takes coordinates of pairs of neurons and outputs a number which is interpreted as the weight of that connection. The advantage this brings is that ANNs can be evolved with complex patterns where collections of neurons have similar behaviour depending on their spatial location. It also means that one evolved function (the CPPN) can determine the strengths of connections of many neurons. It is a form of non-temporal development, where geometrical relationships are translated into weights.

Developmental Symbolic Encoding (DSE) [71] combines concepts from two earlier developmental encodings, Gruau's cellular encoding and L-systems. Like HyperNEAT it can specify connectivity of neurons via evolved geometric patterns. It was shown to outperform HyperNEAT on a shape recognition problem defined over small pixel arrays. It could also produce partly general solutions to a series of even-parity problems of various sizes. Huizinga et al. added an additional output to the CPP program in HyperNEAT that controlled whether or not a connection between a pair of neurons was expressed or not [28]. They showed that the new approach produced more modular solutions and superior performance to HyperNEAT on three specially devised modular problems.

Evolvable-substrate HyperNEAT (ES-HyperNEAT) implicitly defined the positions of the neurons [56], however it proved to be computationally expensive. Iterated ES-HyperNEAT proposed a more efficient way to discover suitable positioning of neurons [58]. This idea was taken further leading to Adaptive HyperNEAT which demonstrated that not only could patterns of weights be evolved but also patterns of local neural learning rules [57]. Like [28] in Adaptive HyperNEAT Risi et al. increased the number of outputs from the CPPN program to encode learning rate and other neural parameters.

3 Abstracting aspects of biological brains

Once we accept that developmental ANNs are desirable, it becomes necessary to abstract and simplify important mechanisms from neuroscience. Which aspects of biological neurons are most relevant depends on the nature of the abstracted neuron model. For instance, if the abstracted developmental neural programs include

genetic regulatory networks one could consider epigenetic processes inside neurons since recent evidence from neuroscience suggests that these may be important in neuron response to past environmental conditions [75, 29]. Here we compare and contrast the proposed abstract model with various aspects of biological neurons in section 3.

Brain development: All multicellular organisms begin as a single cell which undergoes development. Some of the cells become neural stem cells which gradually differentiate into mature neurons. In the proposed model we allow the user to choose how many *non-output* neurons to start with prior to development. However, the model assumes that there is a dedicated output neuron corresponding to each output in the suite of computational problems being solved. Further discussion on the topic of how to handle outputs can be found in Sect. 13.

Arrangement of neurons: The overall architecture of both ANNs and many neural developmental systems is fixed once developed, whereas biological neurons move themselves (in early development) and their branches change over time. Thus the morphology of the network is time dependent and can change during problem solving. In our model, we evolve a developmental process which means that the network of neurons and branches are time dependent.

Neuron structure: Biological neurons have dendrites and axons with branches. Each neuron has a single axon with a variable number of axon branches. In addition, it has a number of dendrites and dendrite branches. There are many types of neurons with different morphologies, some with few dendrites and others with huge numbers of highly branched dendritic trees. In addition neurons have a membrane, a nucleus and a cytoplasm. Since our model of the neuron has zero volume, these aspects are also not included. In our model, the user can choose the minimum and the maximum number of dendrites neurons can have. The evolved developmental programs determine how many dendrites individual neurons can have and indeed, different neurons can have different numbers of dendrites. We have not modelled the axon in our approach. We decided this to keep the proposed model as simple as possible.

Neuron volume: Neurons have many physical properties (volume, temperature, pressure, elasticity...). We have not modelled any of these and like conventional ANNs the neurons in our model are mathematical points. Dendrites are equally unphysical and are merely lines that emanate from neurons and are positioned on the left of the neuron position. They can pass through each other.

Neuron movement: In brains undifferentiated neurons migrate to specific locations in the early stages of development. When they reach their destinations they either develop into mature neurons complete with dendrites and an axon, or they undergo cell death [73]. Moreover, this ability of neurons to migrate to their appropriate positions in the developing brain is critical to brain architecture and func-

tion [46, 54]. We have allowed neuron movement in our model. Neuron movement in real brains is closely related to the physicality of neurons and this could have important consequences. Clearly, mature biological neurons that are entangled with many other neurons will have restricted movement.

Synapses: The connections between biological neurons are called synapses. Signal propagation is via the release of neurotransmitters from the pre-synaptic neuron to the post-synaptic. Like traditional ANNs, we have no notion of neurotransmitters and signals are propagated as if by wires. However, unlike traditional ANNs, dendrites connect with their nearest neuron (on the left). These connections can change when dendrites grow or shrink and when neurons move. Thus, the number of connections between neurons is time-dependent.

Activity Dependent Morphology: There are few proposed models in which changes in levels of activity (in potentials or signals) between neurons leads to changes in neural morphology. This is an extremely important aspect of real brains [53]. This has not been included in the current model. Possibly a measure of signals could be used as an input to the developmental programs. We return to this issue in Sect. 13.

Neuron State: Biological neurons have dendritic trees which change over time. New dendrites can emerge and dendrite branches can develop or be pruned away. Indeed, we discussed in the introduction how important this aspect is for learning. We have abstracted this process by allowing neurons to have multiple dendrites which can replicate (an abstraction of branching) or be removed (an abstraction of pruning). In the model, this process is dependent on a variable called 'health' which is an abstraction of neuron and dendrite state. Khan et al. [35] first suggested the use of this variable.

4 The neuron model

The model we propose is new and conceptually simple. Two evolved neural programs are required to construct neural networks. One represents the neuron soma and the other the dendrite. The role of the soma program is to allow neurons to move, change, die or replicate. For the dendrite, the program needs to be able to grow and change dendrites, cause them to be removed and also to replicate.

The approach is simple in two ways. Firstly, because only two evolved programs are required to build an entire neural network. Secondly, because a snapshot of the neural network at a particular time would show merely a conventional graphs of neurons, weighted connections and standard activation functions. To construct such a developmental model of an artificial neural network we need neural programs that not only apply a weighted sum of inputs to an activation function to determine

the output from the neuron, but a program that can adjust weights, create or prune connections, and create or delete neurons.

Since developmental programs build networks that change over time it is necessary to define new problem classes that are suitable for evaluating such approaches. We argue that trying to solve *multiple* computational problems (potentially even of different types) is an appropriate class of problems.

The pair of evolved programs can be used to build a network from which multiple conventional ANNs can be extracted each of which can solve a different classification problem. We investigate many parameters and algorithmic variants and assess experimentally which aspects are most associated with good performance. Although we have concentrated in this paper on classification problems, our approach is quite general and it could be applied to a much wider variety of problems.

The model is illustrated in in Fig. 1. The neural programs are represented using Cartesian Genetic Programming (CGP) (see Sect. 5). The programs are actually sets of mathematical equations that read variables associated with neurons and dendrites to output updates of those variables. This approach was inspired by some aspects of a developmental method for evolving graphs and circuits proposed by Miller and Thomson [51] and is also strongly influenced by some of the ideas described in [35]. In the proposed model, weights are determined from a program that is a function of neuron position, together with the health, weight and length of dendrites. It is neurocentric and temporal in nature.

As shown in Fig. 1 the inputs to the soma program are: the health, bias and position of the neuron and the average health, length and weight of all dendrites connected to the neuron and problem type. The problem type is a constant (in range [-1, 1]) which indicates whether a neuron is not an output or in the case of an output neuron what computational problem the output neuron belongs to. Let P_t denote the computational problem. Define $P_t = 0$ to denote a non-output neuron, and P_t =1,2 or N_p to respectively denote output neurons belonging to different computational problems, where, N_p denotes the number of computational problems. We define the problem type input to be given by $-1+2P_t/N_p$. For example, if the neuron is not an output neuron the problem type input is -1.0. If it is an output neuron belonging to the last problem its value is 1.0. For all other computational problems its value is a value greater than -1.0 and less than 1.0. The thinking behind the problem type input is that since output neurons are dedicated to a particular computational problem, they should be given information that relates to this, so that the identical neural programs can behave differently according to the computational problem they are associated with. Later experiments were conducted to investigate the utility of problem type (see Sect. 7).

Bias refers to an input to the neuron activation function which is added to the weighted sum of inputs (i.e. it is unweighted). The soma program updates its own health, bias and position based on these inputs. These are indicated by primed symbols in Fig. 1). The user can decide between three different ways of using the program outputs to update the neural variables. The update method is decided by a user defined parameter called Incr_{opt} (see Sec. 4.5) which defines how neuron variables

Fig. 1 The model of a developmental neuron. Each neuron has a position, health and bias and a variable number of dendrites. Each dendrite has a position, health and weight. The behaviour of a neuron soma is governed by a single evolved program. In addition each dendrite is governed by another single evolved program. The soma program decides the values of new soma variables position, health and bias based on previous values, the average over all dendrites belonging to the neuron of dendrite health, position and weight and an external input called *problem type*. The latter is a floating point value that indicates the neuron type. The dendrite program updates dendrite health, position and weight based on previous values, the health, position and bias of the neuron the dendrite belongs to, and the problem type. When the evolved programs are executed, neurons can change, die replicate and grow more dendrites and their dendrites can also change or be removed.

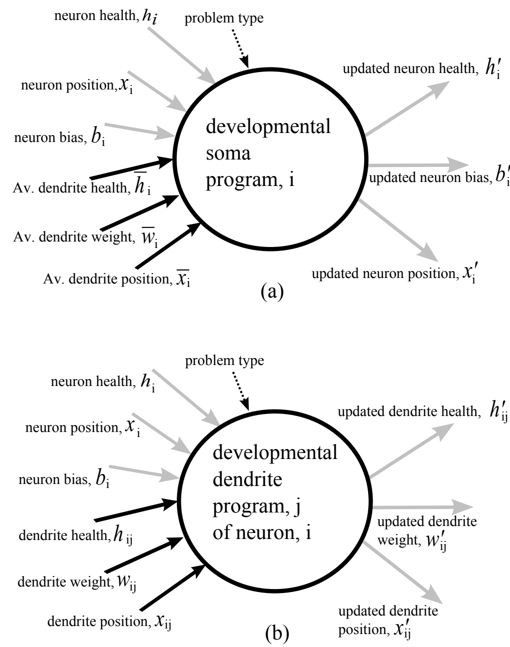

are adjusted by the evolved programs (using user-defined incremental constants or otherwise).

Every dendrite belonging to each neuron is controlled by an evolved dendrite program. As shown in Fig. 1 the inputs to this program are the health, weight and position of the dendrite and also the health, bias and position of the parent neuron. In addition as mentioned earlier, dendrite programs can also receive the problem type of the parent neuron The evolved dendrite program decides how the health, weight and position of the dendrite are to be updated.

In the model, all the neuron and dendrite parameters (weights, bias, health, position and problem type) are defined by numbers in the range $[-1,1]$.

4.1 Fictitious example of early brain development

A fictitious developmental example is shown in Fig. 2. The initial state of the example brain is represented in (a). Initially there is one non-output neuron with a single dendrite. The curved nature of the dendrites is purely for visualisation. In reality the dendrites are horizontal lines emanating from the centre of neurons and of various lengths. When extracting ANNs the dendrites are assumed to connect to their nearest neuron on the left (referred to as ‘snapping’). Output neurons are only allowed to connect to non-output neurons or the first input (by default, if their dendrites lie on the left of the leftmost non-output neuron). Thus the ANN that can be extracted from the initial example brain, has three neurons. The non-output neuron is connected to the second input and both output neurons are connected via their single dendrite to the non-output neuron.

Fig. 2 Example showing a developing fictitious example brain. The squares on the left represent the inputs. The solid circles indicate non-output neurons. Non-output neurons have solid dendrites. The dotted circles represent output neurons. Output neuron’s dendrites are also dotted. In this example, we assume that only output neurons are allowed to move. The neurons, inputs and dendrites are all bound to the interval [-1,1]. Dendrites connect to nearest neurons or inputs on the *left* of their position (*snapping*). (a) shows the initial state of the example brain. (b) shows the example brain after one developmental step and (c) shows it after two developmental steps.

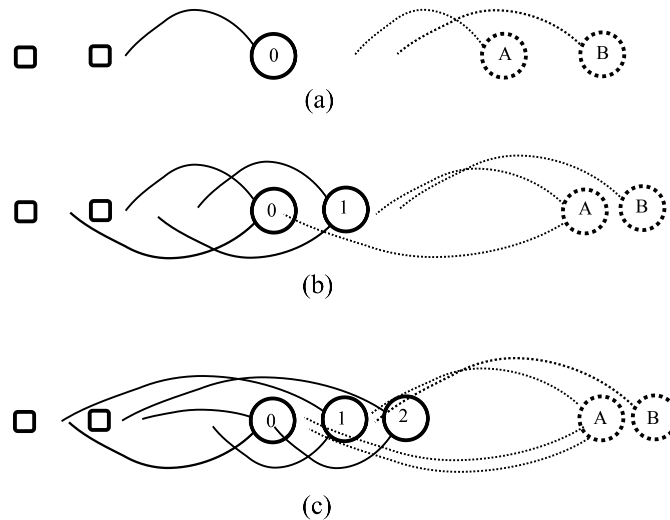

Fig. 2(b) shows the example brain after a single developmental step. In this step, the soma program and dendrite programs are executed in each neuron. The non-output neuron (labeled 0) has replicated to produce another non-output neuron (labeled 1) it has also grown a new dendrite. Its dendrites connect to both inputs. The

newly created non-output neuron is identical to its parent except that its position is a user-defined amount, MN_{inc}, to the right of the parent and its health is set to 1 (an assumption of the model). Both its dendrites connect to the second input. It is assumed that the soma programs running in the two output neurons A and B have resulted in both output neurons having moved to the right. Their dendrites have also grown in length. Neuron A's first dendrite is now connected to neuron one. In addition, neuron A has high health so that it has grown a new dendrite. Every time a new dendrite grows it is given a weight and health equal to 1.0. Also its new dendrite is given a position equal to half the parent neuron's position. These are assumptions of the model. Thus its new dendrite is connected to neuron zero. Neuron B's only dendrite is connected to neuron one.

Fig. 2(c) shows the example brain after a two developmental steps. The dendrites of neuron zero have changed little and it is still connected in the same way as the previous step. The dendrites of neuron one have both changed. The first one has become longer but remains connected to the first input. The second dendrite has become shorter but it still snaps to the second input. Neuron one has also replicated as a result of its health being above the replication threshold. It gets dendrites identical to its parent, its position is again incremented to the right of its parent and its health is set to 1.0. Its first dendrite connects to input one and its second dendrite to neuron zero. Output neuron A has gained a dendrite, due to its health being above the dendrite birth threshold. The new dendrite stretches to a position equal to half of its parent neuron. So it connects to neuron zero. The other two dendrites remain the same and they connect to neuron one and zero respectively. Finally, output neuron B's only dendrite has extended a little but still snaps to neuron one. Note, that at this stage neuron two is not connected to by another neuron and is redundant. It will be stripped out of the ANN that is extracted from the example brain.

4.2 Model parameters

The model necessarily has a large number of user-defined parameters these are shown in Table 1.

The total number of neurons allowed in the network is bounded between a user-defined lower (upper)bound NN_{min} (NN_{max}). The number of dendrites each neuron can have is bounded by user-defined lower (upper) bounds denoted by DN_{min} (DN_{max}). These parameters ensure that the number of neurons and connections per neuron remain in well-defined bounds, so that a network can not eliminate itself or grow too large. The initial number of neurons is defined by N_{init} and the initial number of dendrites per neuron is given by ND_{init}.

If the health of a neuron falls below (exceeds) a user-defined threshold, NH_{death} (NH_{birth}) the neuron will be deleted (replicated). Likewise, dendrites are subject to user defined health thresholds, DH_{death} (DH_{birth}) which determine whether the dendrite will be deleted or a new one will be created. Actually, to determine dendrite birth the parent *neuron* health is compared with DH_{birth} rather than dendrite

health. This choice was made to prevent the potential very rapid growth of dendrite numbers.

When the soma or dendrite programs are run the outputs are used to decide how to adjust the neural and dendrite variables. The amount of the adjustments are decided by the six user-defined δ parameters.

The number of developmental steps in the two developmental phases ('pre' learning and 'while' learning) are defined by the parameters, NDS_{pre} and NDS_{whi}. The number of learning epochs is defined by N_{ep}. Note that the pre-learning phase of development, 'pre', can have different incremental constants (i.e. δs) to the learning epoch phase, 'while'.

In some cases, neurons will collide with other neurons (by occupying a small interval around an existing neuron) and the neuron has to be moved by a certain increment until no more collisions take place. This increment is given by MN_{inc}.

The places where external inputs are provided is predetermined uniformly within the region between -1 and I_u. The parameter I_u defines the upper bound of their position. Also output neurons are initially uniformly distributed between the parameter O_l and 1. However, depending on a user-defined option the output neurons as with other neurons can move according to the neuron program. All neurons are marked as to whether they provide an external output or not. In the initial network there are N_{init} non-output neurons and N_o output neurons, where N_o denotes the number of outputs required by the computational problem being solved.

Finally, the neural activation function (hyperbolic tangent) and the sigmoid function (which is used in nonlinear incremental adjustment of neural variables) have a slope constant given by α.

Table 1 Table of neural model constants and their meanings.

Symbol	Meaning
$NN_{min}(NN_{max})$	Min. (Max.) allowed number of neurons
N_{init}	Initial number of non-output neurons
$DN_{min}(DN_{max})$	Min. (Max.) number of dendrites per neuron
ND_{init}	Initial number of dendrites per neuron
$NH_{death}(NH_{birth})$	Neuron health thresholds for death (birth)
$DH_{death}(DH_{birth})$	Dendrite health thresholds for death (birth)
δ_{sh}	Soma health increment (pre, while)
δ_{sp}	Soma position increment (pre, while)
δ_{sb}	Soma bias increment (pre, while)
δ_{dh}	Dendrite health increment (pre, while)
δ_{dp}	Dendrite position increment (pre, while)
δ_{dw}	Dendrite weight increment (pre, while)
NDS_{pre}	Number of developmental steps before epoch
NDS_{whi}	Number of 'while' developmental steps during epoch
N_{ep}	Number of learning epochs
MN_{inc}	Move neuron increment if collision
I_u	Max. program input position
O_l	Min. program output position
α	Sigmoid/Hyperbolic tangent exponent constant

4.3 Developing the brain and evaluating the fitness

An overview of the algorithm used for training and developing the ANNs is given in Overview 1.

Overview 1 Overview of fitness algorithm

```
 1: function FITNESS
 2:     Initialise brain
 3:     Load 'pre' development parameters
 4:     Update brain NDS_pre times by running soma and dendrite programs
 5:     Load 'while' developmental parameters
 6:     repeat
 7:         Update brain NDS_whi times by running soma and dendrite programs
 8:         Extract ANN for each benchmark problem
 9:         Apply training inputs and calculate accuracy for each problem
10:         Fitness is the normalised average accuracy over problems
11:         If fitness reduces terminate learning loop and return previous fitness
12:     until N_ep epochs complete
13:     return fitness
14: end function
```

The brain is always initialised with at least as many neurons as the maximum number of outputs over all computational problems. Note, all problem outputs are represented by a unique neuron dedicated to the particular output. However, the maximum and initial number of non-output neurons can be chosen by the user. Non-output neurons can grow change or give birth to new dendrites. Output neurons can change but not die or replicate as the number of output neurons is fixed by the choice of computational problems. The detailed algorithm for training and developing the ANN is given in Algorithm 1.

Development of the brain can happen in two phases, 'pre' and 'while'. The 'pre' phase runs for NDS_{pre} developmental steps and is outside the learning loop. This is an early phase of development. It has its own set of developmental parameters. The 'while' phase happens inside the learning loop which has N_{ep} epochs. It too has its own set of developmental parameters. The idea of two phases is inspired by the phases of biological development. In early brain development, neurons are stem cells that move to particular locations and have no dendrites and axons. This can effectively be mimicked since in the 'pre' phase the parameters controlling dendrites can be disabled by setting $DH_{death} = -1.0$ and $DH_{birth} = 1.0$. This means that dendrites can not be removed or be born. In addition, setting $\delta_{dh} = 0.0$, $\delta_{dp} = 0.0$ and $\delta_{dw} = 0.0$ would mean that any existing dendrites could not change. In a similar way, 'while' parameters could be chosen to disallow somas to move, die, replicate or change during the learning loop and also to allow dendrites to grow/shrink, change, be removed, or replicate. Thus it can be seen that the collection of parameters gives the user a lot of control of the developmental process.

The learning loop evaluates the brain by extracting conventional ANNs for each problem and calculating a fitness (based on accuracy of classification) it checks to see if the new fitness value is greater than or equal to the previous value at the last epoch. If the fitness has reduced the learning loop terminates and the previous value of fitness is returned. The purpose of the learning loop is to enable the evolution of a development process that progressively improves the brain's performance. The aim is to find programs for the soma and dendrite which allow this improvement to continue with epoch beyond the limit chosen (N_{ep}). Later in this chapter, an experiment is conducted to test whether this general learning behaviour has been achieved (see Sect. 12).

4.4 Updating the brain

Updating the brain is the process of running the soma and dendrite programs once in all neurons and dendrites (i.e. it is a single developmental step). Doing this will cause the brain to change and after all changes have been carried out a new updated brain will be produced. This replaces the previous brain. Overview algorithm 2 gives a high-level overview of the update brain process.

Overview 2 Update brain overview

```
1:  function UPDATEBRAIN
2:      Run soma program in non-output neurons to update soma
3:      Ensure neuron does not collide with neuron in updated brain
4:      Run dendrite program in all non-output neurons
5:      If neuron survives add it to updated brain
6:      If neuron replicates ensure new neuron does not collide
7:      Add new neuron to updated brain
8:      Run soma program in output neurons to update soma
9:      Ensure neuron does not collide
10:     Run dendrite program in all output neurons
11:     If neuron survives add it to updated brain
12:     Replace old brain with updated brain
13: end function
```

Sect. 15.1 presents a more detailed version of how the brain is updated at each developmental step (see Algorithm 2) and gives details of the neuron collision avoidance algorithm.

4.5 Running and updating the soma

The UPDATEBRAIN program calls the RUNSOMA program to determine how the soma changes in each developmental step. As we saw in Fig. 1(a) the seven soma program inputs are: neuron health, position and bias, the averaged position, weight and health of the neuron's dendrites and the problem type. Once the evolved CGP soma program is run the soma outputs are returned to the brain update program. These steps are shown in Overview 2.

Overview 2 Running the soma: algorithm overview

```
1: function RUNSOMA
2:     Calculate average dendrite health, position and weight
3:     Gather soma program inputs
4:     Run soma program
5:     Return updated soma heath, bias and position
6: end function
```

The detailed version of the RUNSOMA function can be found in Sect. 15.3. The RUNSOMA function uses the soma program outputs to adjust the health, position and bias of the soma according to three user-chosen options defined by a variable Incr_{opt}.This is carried out by the UPDATENEURON overview Alg. 3.

Overview 3 Update neuron algorithm overview

```
1: function UPDATENEURON
2:     Assign original neuron variables to parent variables
3:     Assign outputs of soma program to health, position and bias
4:     Depending on Incr_opt get increments
5:     If soma program outputs > 0 (≤0) then incr. (decr.) parent variables
6:     Assign parent variables to neuron
7:     Bound health, position and bias
8: end function
```

4.6 Updating the dendrites and building the new neuron

This section is concerned with running the evolved dendrite programs. In every dendrite, the inputs to the dendrite program have to be gathered. The dendrite program is executed and the outputs are used to update the dendrite. This is carried out by a function called RUNDENDRITE. Note, in RUNALLDENDRITES we build the completely updated neuron from the updated soma and dendrite variables. The sim-

plified algorithm for doing this is shown in overview algorith 4. The more detailed version is available in Sect. 15.5.

Overview 4 An overview of the RUNALLDENDRITES algorithm which runs all dendrite programs and uses all updated variables to build a new neuron.

```
1: function RUNALLDENDRITES
2:     Write updated soma variables to new neuron
3:     if Old soma health > DH_birth then
4:         Generate a dendrite for new neuron
5:     end if
6:     for all Dendrites do
7:         Gather dendrite program inputs
8:         Run dendrite program to get updated dendrite variables
9:         Run dendrite to get updated dendrite
10:        if Updated dendrite is alive then
11:            Add updated dendrite to new neuron
12:            if Maximum number of dendrites reached then
13:                Stop processing any more dendrites
14:            end if
15:        end if
16:    end for
17:    if All dendrites have died then
18:        Give new neuron the first dendrite of the old neuron
19:    end if
20: end function
```

Overview Alg. 4 (in line 9) uses the updated dendrite variables obtained from running the evolved dendrite program to adjust the dendrite variables (according to the incrementation option chosen). This function is shown in the overview Alg. 5. The more detailed version is available in Sect. 15.5.

The RUNDENDRITE function begins by assigning the dendrite's health, position and weight to the parent dendrite variables. It writes the dendrite program outputs to the internal variables health, weight and position. It respectively carries out the increments or decrements of the parent dendrite variables according whether the corresponding dendrite program outputs are greater than or less than or equal to zero. After this it bounds those variables. Finally, it updates the dendrites health, weight and position provided the adjusted health is above the dendrite death threshold.

We saw in the fitness function that we extract conventional ANNs from the evolved brain. The way this is accomplished is as follows.

Since we share inputs across problems we set the number of inputs to be the maximum number of inputs that occur in the computational problem suite. If any problem has less inputs the extra inputs are set to zero.

The next phase is to go through all dendrites of the neurons to determine which inputs or neurons they connect to. To generate a valid neural network we assume that dendrites are automatically connected to the nearest neuron or input on the left. We refer to this as *snapping*. The dendrites of non-output neurons are allowed to connect

Overview 5 Change dendrites according to the evolved dendrite program.

```
1: function RunDendrite
2:     Assign original dendrite variables to parent variables
3:     Assign outputs of dendrite program to health, position and weight
4:     Depending on Incr_opt get increments
5:     If dendrite program outputs > 0 (≤0) then incr(decr.) parent variables
6:     Assign parent variables to neuron
7:     Bound health, position and weight
8:     if (health > DH_death) then
9:         Update dendrite variables
10:        Dendrite is alive
11:    else
12:        Dendrite is dead
13:    end if
14:    Return updated dendrite variables and whether dendrite is alive
15: end function
```

to either inputs or other non-output neurons on their left. However, output neurons are only allowed to connect to non-output neurons on their left. It is not desirable for the dendrites of output neurons to be connected directly to inputs, however, when output neurons are allowed to move, they may only have inputs on their left. In this case the output neuron's dendrite will be connected to the first external input to the ANN network (by default).

The detailed version of the ANN extraction process is given in Sect. 15.6.

5 Cartesian GP

The two neural programs are represented and evolved using a form of Genetic Programming (GP) known as Cartesian Genetic Programming (CGP). CGP [48, 50] is a form of GP in which computational structures are represented as directed, often acyclic graphs indexed by their Cartesian coordinates. Each node may take its inputs from any previous node or program input (although recurrent graphs can also be implemented see [77]). The program outputs are taken from the output of any internal node or program input. In practice, many of the nodes described by the CGP chromosome are not involved in the chain of connections from program input to program output. Thus, they do not contribute to the final operation of the encoded program, these inactive, or "junk", nodes have been shown to greatly aid the evolutionary search [49, 79, 81]. The representational feature of inactive genes in CGP is also closely related to the fact that it does not suffer from bloat [47].

In general, the nodes described by CGP chromosomes are arranged in a rectangular $r \times c$ grid of nodes, where r and c respectively denote the user-defined number of rows and columns. In CGP, nodes in the same column are not allowed to be connected together. CGP also has a connectivity parameter l called "levels-back"

which determines whether a node in a particular column can connect to a node in a previous column. For instance if $l = 1$ all nodes in a column can only connect to nodes in the previous column. Note that levels-back only restricts the connectivity of nodes; it does not restrict whether nodes can be connected to program inputs (terminals). However, it is quite common to adopt a linear CGP configuration in which $r = 1$ and $l = c$. This was done in our investigations here. CGP chromosomes can describe multiple input multiple output (MIMO) programs with a range of node functions and arities. For a detailed description of CGP, including its current developments and applications, see [48]. Both the soma and dendrite program have 7 inputs and 3 outputs. (see Fig. 1). The function set chosen for this study are defined over the real-valued interval [-1.0, 1.0]. Each primitive function takes up to three inputs, denoted z_0, z_1 and z_2. The functions are defined in Table 2.

Table 2 Node function gene values, mnemonic and function definition

Value	mnemonic	Definition
0	abs	$\|z_0\|$
1	sqrt	$\sqrt{\|z_0\|}$
2	sqr	$z_0{}^2$
3	cube	$z_0{}^3$
4	exp	$(2exp(z_0+1)-e^2-1)/(e^2-1)$
5	sin	$\sin(z_0)$
6	cos	$\cos(z_0)$
7	tanh	$\tanh(z_0)$
7	inv	$-z_0$
9	step	**if** $z_0 < 0.0$ **then** 0 **else** 1.0
10	hyp	$\sqrt{(z_0{}^2+z_1{}^2)/2}$
11	add	$(z_0+z_1)/2$
12	sub	$(z_0-z_1)/2$
13	mult	$z_0 z_1$
14	max	**if** $z_0 >= z_1$ **then** z_0 **else** z_1
15	min	**if** $z_0 <= z_1$ **then** z_0 **else** z_1
16	and	**if** $(z_0 > 0.0$ **and** $z_1 > 0.0)$ **then** 1.0 **else** -1.0
17	or	**if** $(z_0 > 0.0$ **or** $z_1 > 0.0)$ **then** 1.0 **else** -1.0
18	rmux	**if** $z_2 > 0.0$ **then** z_0 **else** z_1
19	imult	$-z_0 z_1$
20	xor	**if** $(z_0 > 0.0$ **and** $z_1 > 0.0)$ **then** -1.0 **else if** $(z_0 < 0.0$ **and** $z_1 < 0.0)$ **then** -1.0 **else** 1.0
21	istep	**if** $z_0 < 1.0$ **then** 0 **else** -1.0
22	tand	**if** $(z_0 > 0.0$ **and** $z_1 > 0.0)$ **then** 1.0 **else if** $(z_0 < 0.0$ **and** $z_1 < 0.0)$ **then** -1.0 **else** 0.0
23	tor	**if** $(z_0 > 0.0$ **or** $z_1 > 0.0)$ **then** 1.0 **else if** $(z_0 < 0.0$ **or** $z_1 < 0.0)$ **then** -1.0 **else** 0.0

6 Benchmark problems

In this study, we evolve neural programs that build ANNs for solving three standard classification problems. The problems are cancer, diabetes and glass. The definitions of these problems are available in the well-known UCI repository of machine learning problems[1]. These three problems were chosen because they are well-studied and also have similar numbers of inputs and a small number of classes. Cancer has 9 real attributes and two Boolean classes. Diabetes has 8 real attributes and two Boolean classes. Glass has 9 real attributes and six Boolean classes. The specific datsets chosen were cancer1.dt, diabetes1.dt and glass1.dt which are described in the PROBEN suite of problems [2].

7 Experiments

The long-term aim of this research is to explore effective ways to *develop* ANNs. The work presented here is just a beginning and there are many aspects that need to be investigated in the future (see Sect. 13). The specific research questions we have focused on in this chapter are:

- What types of neuron activation function is most effective?
- How many neurons and dendrites should we allow?
- Should neuron and dendrite programs be allowed to read problem type?

These questions complement the questions asked and investigated using the same neural model in recent previous work [52]. There, the utility of neuron movement in both non-output and output neurons was investigated. It was found that statistically significantly better results were obtained when only output-neurons were allowed to move. In addition, the work examined three ways of incrementing or decrementing neural variables. In the first the outputs of evolved programs determines directly the new values of neural variables (position, health, bias, weight), that is to say there is no incremental adjustment of neural variables. In the second, the variables are incremented or decremented in user-defined amounts (the deltas in Table 1). In the third, the adjustments to the neural variables are nonlinear (they are adjusted using a sigmoid function). Linear adjustment of variables (increment or decrement) was found to be statistically superior to the alternatives.

To answer the questions above, a series of experiments were carried out to investigate the impact of various aspects of the neural model on classification accuracy. Twenty evolutionary runs of 20,000 generations of a 1+5-ES were used. Genotype lengths for soma and dendrite programs were chosen to be 800 nodes. Goldman mutation [20, 21] was used which carries out random point mutation until an active gene is changed. For these experiments a subset of allowed node functions were

[1] `https://archive.ics.uci.edu/ml/datasets.html`

[2] `https://publikationen.bibliothek.kit.edu`

chosen as they appeared to give better performance. These were: step, add, sub, mult, xor, istep. The remaining experimental parameters are shown in Table 3:

Some of the parameter values are very precise (defined to the fourth decimal place). The process of discovering these values consisted of an informal but greedy search for good parameters that produced high fitness in the first evolutionary run. It was fortuitous as it turned out that this was a reasonable way of obtaining good parameters *on average*.

Table 3 Table of neural model parameters.

Parameter	Value
$NN_{min}(NN_{max})$	0 (20-100)
N_{init}	5
$DN_{min}(DN_{max})$	1 (5-50)
ND_{init}	5
NDS_{pre}	4
NDS_{whi}	8
N_{ep}	1
MN_{inc}	0.03
I_u	-0.6
O_l	0.8
α	1.5
'Pre' development parameters	
$NH_{death}(NH_{birth})$	-0.405 (0.406)
$DH_{death}(DH_{birth})$	-0.39 (-0.197)
δ_{sh}	0.1
δ_{sp}	0.1487
δ_{sb}	0.07
δ_{dh}	0.1
δ_{dp}	0.1
δ_{dw}	0.101
'While' development parameters	
$NH_{death}(NH_{birth})$	0.435 (0.7656)
$DH_{death}(DH_{birth})$	0.348 (0.41)
δ_{sh}	0.009968
δ_{sp}	0.01969
δ_{sb}	0.01048
δ_{dh}	0.0107
δ_{dp}	0.0097
δ_{dw}	0.0097

8 Results

The mean, median, maximum and minimum accuracies achieved over 20 evolutionary runs for three different neuron activation functions neurons are shown in Table 4. We can see that the best values of mean, median, maximum and minimum

are all obtained when the hyperbolic tangent function is used. The rectilinear neural functions is zero for negative arguments and equal to its argument for positive arguments. Both the sigmoid and hyperbolic tangent activation functions have α as exponent multipliers.

Table 4 Training and testing accuracy for three neural activation functions.

Acc.	Hyperbolic Tangent Train (Test)	Rectilinear Train (Test)	Sigmoid Train (Test)
Mean	0.7401 (0.7075)	0.7150 (0.6698)	0.6980 (0.6760)
Median	0.7392 (0.7266)	0.7101 (0.6717)	0.7036 (0.6851)
Maximum	0.7988 (0.7669)	0.7654 (0.7398)	0.7315 (0.7237)
Minimum	0.6840 (0.6200)	0.6815 (0.6217)	0.6251 (0.5894)

Table 5 Training and testing accuracy on individual problems when using tanh activation.

Acc.	Cancer Train (Test)	Diabetes Train (Test)	Glass Train (Test)
Mean	0.9086 (0.9215)	0.7233 (0.6594)	0.5883 (0.5415)
Median	0.9129 (0.9281)	0.7266 (0.6615)	0.5841 (0.5849)
Maximum	0.9657 (0.9770)	0.7630 (0.6823)	0.6729 (0.6981)
Minimum	0.8571 (0.8621)	0.6693 (0.6198)	0.4673 (0.3396)

Table 6 shows how the results for the model (using tanh activation) compare with the performance of 179 classifiers (covering 17 families) [15][3]. The figures are given just to show that the results for the developmental ANNs are respectable. The results are particularly encouraging considering that the evolved developmental programs build classifiers for three different classification problems simultaneously, so the comparison is unfairly biased against the proposed model. The cancer results produced by the model are very close to those compiled from the suite of ML methods, however it can be seen that the models's results for diabetes and glass are not as close. It is unclear why this is the case.

Table 6 Comparison of test accuracies on three classification problems. Model using tanh activation compared with huge suite of classification methods as described in [15]

Acc.	Cancer ML (model)	Diabetes ML (model)	Glass ML (model)
Mean	0.935(0.9215)	0.743(0.6594)	0.610(0.5415)
Maximum	0.974(0.9770)	0.790(0.6822)	0.785(0.6981)
Minimum	0.655(0.8621)	0.582(0.6198)	0.319(0.3340)

[3] The paper gives a link to the detailed performance of the 179 classifiers which contain the figures given in the table

The second series of experiments examined how the classifier performance varied with the upper bound on the number of allowed neurons and the number of dendrites.

Table 7 Training and testing accuracy for different upper bounds on the number of neurons (NN_{max}). The number of dendrites was 50 in all cases.

Acc.	NN_{max}=20 Train (Test)	NN_{max}=40 Train (Test)	NN_{max}=60 Train (Test)	NN_{max}=80 Train (Test)	NN_{max}=100 Train (Test)
Mean	0.7314 (0.7026)	**0.7401** (0.7075)	0.7161 (0.6927)	0.7179 (0.6969)	0.7222 (0.7072)
Median	0.7327 (0.7176)	**0.7392** (0.7266)	0.7155 (0.6919)	0.7252 (0.7045)	0.7219 (0.7081)
Maximum	0.7605 (0.7468)	**0.7988** (0.7669)	0.7355 (0.7461)	0.7757 (0.7483)	0.7493 (0.7480)
Minimum	**0.7002** (0.6563)	0.6840 (0.6200)	0.6654 (0.6408)	0.6284 (0.5998)	0.6691 (0.6492)

The third series of experiments was concerned with varying the upper bound on the number of dendrites allowed for each neuron (DN_{max}). It was found that the results when the dendrite upper bound is 20 produced exactly the same results as an upper bound of 50. All dendrites must snap to a neuron or input and therefore contribute to the output of the network. The fact that increasing the upper bound made no difference implies that the number of dendrites the neurons used never exceeded 20 anyway.

Table 8 Training and testing accuracy for different upper bounds on the number of dendrites (DN_{max}). The upper bound on the number of neurons was 40 in all cases.

Acc.	DN_{max}=5 Train (Test)	DN_{max}=10 Train (Test)	DN_{max}=15 Train (Test)	DN_{max}=50 Train (Test)
Mean	**0.7408** (0.7125)	0.7347 (0.7053)	0.7338 (0.6966)	0.7401 (0.7075)
Median	**0.7463** (0.7208)	0.7335 (0.7152)	0.7325 (0.6857)	0.7392 (0.7266)
Maximum	0.7926 (**0.7701**)	0.7924 (0.7640)	0.7988 (0.7669)	**0.7988** (0.7669)
Minimum	0.6789 (0.6047)	0.6850 (0.6395)	0.6797 (0.6281)	**0.6840** (**0.6200**)

Statistical significance testing with the test data revealed that the better scenarios (NN_{max}=20, 40 or DN_{max}=5, 50) are only weakly statistically different from other scenarios. So it appears that performance is not particularly sensitive to the maximum number of neurons and dendrites (within reasonable bounds).

The fourth series of experiments concerned the issue of problem type. As discussed in Section 4, problem type is a real-valued quantity in the interval [-1, 1]. It is a quantity that indicates what computational problem a neuron belongs to. Non-output neurons are not committed to a problem type so the problem type is assumed to be -1.0. However, output neurons are all dedicated to a particular problem. The cancer problem has two output neurons, the diabetes has two and the glass problem has six. The extraction process that gets the ANNs associated with each problem begins from the output neurons corresponding to each problem. So the question arises as to whether it is useful or not to allow the neural programs to read the problem type? The results are shown in Table 9.

Table 9 Training and testing accuracy with and without problem type inputs to evolved programs. The upper bounds on the number of neurons and dendrites was 40 and 50 respectively.

Acc.	Without problem type Train (Test)	With problem type Train (Test)
Mean	0.7314 (0.7026)	**0.7401** (**0.7075**)
Median	0.7327 (0.7176)	**0.7392** (**0.7266**)
Maximum	0.7605 (0.7468)	**0.7988** (**0.7669**)
Minimum	**0.7002** (**0.6563**)	0.6840 (0.6200)

Using the problem type as an input to neural programs appears to be useful as it improves the mean, median and maximum compared with not using the problem type. However, tests of statistical significance discussed in the next section show that the difference is not significant.

9 Comparisons and statistical significance

The Wilcoxon Ranked-Sum test (WRS) was used to assess the statistical difference between pairs of experiments. In this test, the null hypothesis is that the results (best accuracy) over the multiple runs for the two different experimental conditions are drawn from the same distribution and have the same median. If there is a statistically significant difference between the two then null hypothesis is false with a degree of certainty which depends on the smallness of a calculated statistic called a p-value. However, in the WRS before interpreting the p-value one needs to calculate another statistic called Wilcoxon's W value. This value needs to be compared with calculated values which depend on the number of samples in each experiment. Results are statistically significant when the calculated W-value is less than or equal to certain critical values for W[82]. The critical values depend on the sample sizes and the p-value. We used a publicly available Excel spreadsheet for doing these calculations[4]. The critical W-values can be calculated in two ways: one-tailed or two-tailed. The two-tailed test is appropriate here as we are interested in whether one experiment is better than another (and vice versa).

For example, in Table 10 the calculated W-value is 34 and the critical W-value for for the paired sample sizes of 20 (number of runs) with p-value less than 0.01 is 38 (assuming a two-tailed test)[5]. The p-value gives a measure of the certainty with which the null hypothesis can be accepted. Thus the lower the value the more likely that the two samples come from different distributions (i.e. are statistically different). Thus in this case, the probability that the null hypothesis can be rejected is 0.99.

[4] http://www.biostathandbook.com/wilcoxonsignedrank.html

[5] http://www.real-statistics.com/statistics-tables/wilcoxon-signed-ranks-table/

Table 10 Statistical comparison of *testing* results from experiments (Wilcoxon Rank-Sum two-tailed).

Question	Expt. A	Expt. B	W	W critical	P-value	significant?
Activation	tanh	rectilinear	34	38	$0.005 < p < 0.01$	yes
Activation	rectilinear	sigmoid	90	69	$0.2 < p$	no
Problem Type Input	Yes	No	92	69	$0.2 < p$	no

The conclusions we can draw from Table 10 are that a hyperbolic tangent activation function is statistically significantly superior to either a rectilinear or sigmoid function. In addition, the use of problem type as an evolved program input is not statistically distinguishable from not using a problem type input. This is surprising as output neurons are dedicated to problem types and one might expect by a problem type input would allow neural programs to behave differently according to problem type.

10 Evolved developmental programs

The average number of *active* nodes in the soma and dendrite programs for the hyperbolic tangent activation function is 54.7 in both cases. Thus the programs are relatively simple. It is also possible that the graphs can be logically reduced to even simpler forms. The graphs of the active nodes in the CGP graphs for the best evolutionary run (0) are shown in Figs. 3 and 4. The red input connections between nodes indicate the first input in the subtraction operation. This is the only node operation where node input order is important.

Fig. 3 Best evolved soma program. The input nodes are: soma heath (sh), soma bias (sb), soma position (sp), average dendrite health (adh), average dendrite weight (adw), average dendrite position (adp) and problem type (pt). The output nodes are: soma health (SH), soma bias (SB) and soma position (SP).

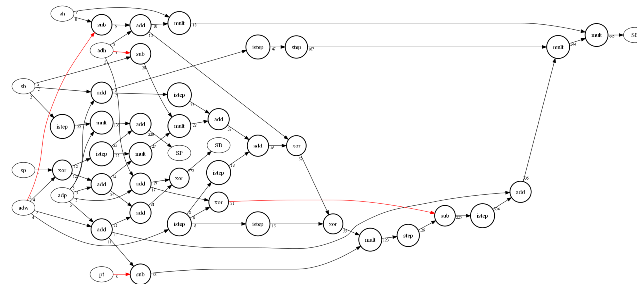

Fig. 4 Best evolved dendrite program. The input nodes are: soma heath (sh), soma bias (sb), soma position (sp), dendrite health (dh), dendrite weight (dw), dendrite position (dp) and problem type (pt). The output nodes are: dendrite health (DH), dendrite weight (DW) and dendrite position (DP).

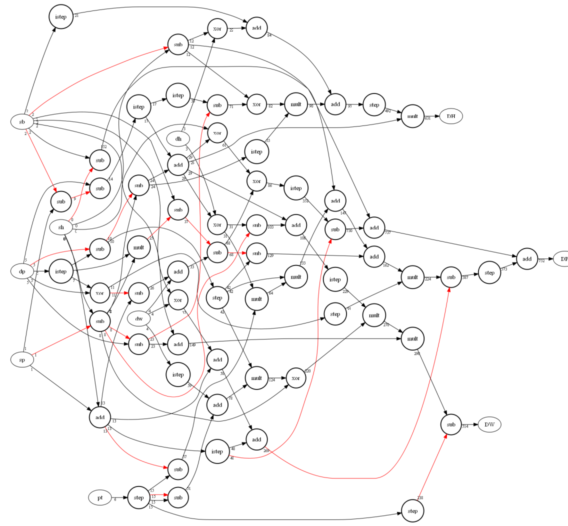

11 Developed ANNs for each classification problem

The ANNs for the best evolutionary run (0) were extracted (using Alg. 9) and can be seen in Figs. 5 and 6. The plots ignore connections with weight equal to zero.

In the figures, a colour scheme is adopted to show which neurons belong to which problems. Red indicates a neuron belonging to the ANN predicting cancer, green indicates a neuron belonging to the ANN predicting diabetes and blue indicates a neuron belonging to an ANN predicting the type of glass. If neurons are shared the colours are a mixture of the primary colours. So white, or an absence of colour, indicates that a neuron is shared over all three problems. Magenta indicates a neurons shared between cancer and glass. Yellow would indicate neurons shared between cancer and diabetes (in the case shown there are no neurons shared between cancer and diabetes).

The ANNs use 21 neurons in total of which only four non-output neurons were unshared (output neurons cannot be shared). It is interesting to note that the number of neurons used for the diabetes data set was 6 (two of which were output neurons) even though the diabetes classification problem is more difficult than cancer. Also the brain was allowed to use up to 40 neurons but many of these were redundant and not used in the extracted ANNs. These redundant neurons were not referenced in the chain of connections from output neurons to inputs. This redundancy is very much like the redundancy that occurs in CGP when extracting active computational nodes.

Fig. 5 Developed ANN for cancer dataset. This dataset has 9 attributes and two outputs. The numbers inside the circles are the neuron bias. The numbers in larger font near but outside the neurons are neuron IDs. The connection weights between neurons are shown near to the connections. If any attributes are not present it means they are unused. The training accuracy is 0.9657 and the test accuracy is 0.9770.

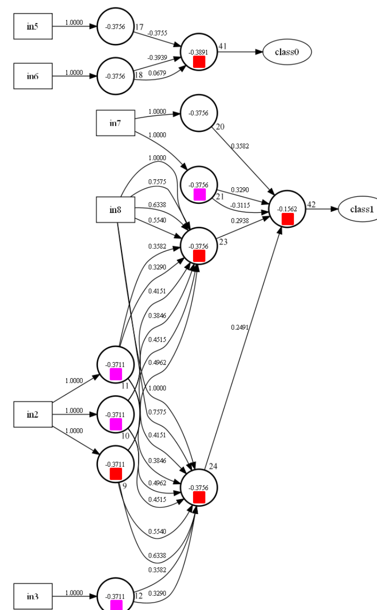

It is interesting that for the cancer prediction ANN, class 0 is decided by a very small network in which only two attributes are involved (the 5th and 6th attributes). This network is independent of the network determining class 1. It is also interesting that the neurons with IDs 17 and 18 are actually identical. Note that attributes 0, 1 and 4 have been ignored. There are three identical neurons connected to attribute 2 (with IDs 9, 10 and 11), these can be reduced to a single neuron!

Fig. 6 Developed ANN for diabetes (a) and glass (b) datasets.The diabetes dataset has 8 attributes and two outputs. The glass dataset has has 9 attributes and six outputs. The numbers inside the circles are the neuron bias. The numbers in larger font near but outside the neurons are neuron IDs. The connection weights between neurons are shown near to the connections. Attributes not present are unused. The diabetes training accuracy is 0.7578 and the test accuracy is 0.6823. The glass training accuracy is 0.6729 and the test accuracy is 0.6415

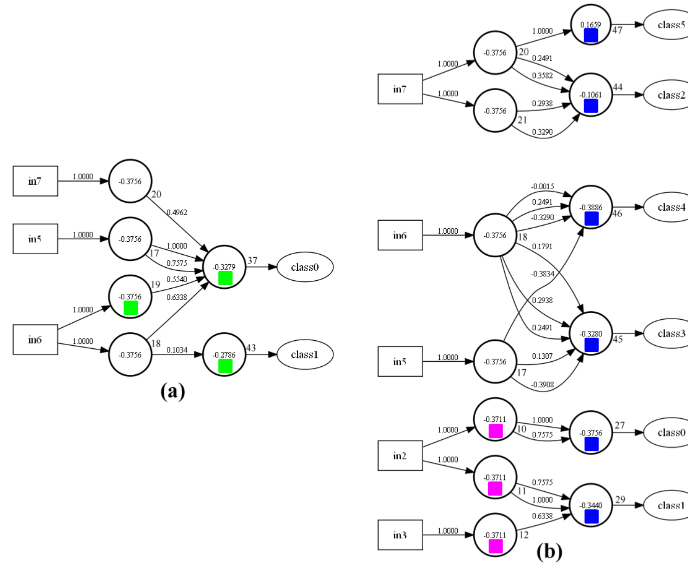

It is surprising that the diabetes ANN is so small. Indeed the neurons with IDs 18 and 19 can be replaced with a single neuron as the two neurons are identical. The ANN only reads three attributes of the dataset! In the case of the glass ANN, once again we see that the ANN consists of distinct networks (3), one predicting classes 2 and 5, one predicting classes 0, 1 and one predicting classes 3 4.

When we analyzed ANNs produced in other cases we sometimes observed ANNs in which neurons occur that only have inputs with weight zero (i.e. effectively no inputs). Such neurons can still be involved in prediction as provided the bias is non-zero the neuron will output a constant value.

Another interesting feature is that pairs of neurons often have multiple connections. This is equivalent to a single connection where the weighted value is the sum of the individual connections weights. This phenomenon was also observed in CGP encoded and evolved ANNs [76].

12 Evolving neural learning programs

The fitness function (see overview algorithm 1) included the possibility of learning epochs. In this section we present and discuss results when a number of learning epochs have been chosen. The task for evolution is then to construct two neural programs that develop ANNs that improve with each learning epoch. The aim is to find a general learning algorithm in which the ANNs change and improve with each learning epoch beyond the limited number of epochs used in training. The 'while' experimental parameters required to investigate were changed from those used previously when there were no learning epochs. The new parameters are shown in Table 11.

Table 11 Table of neural model parameters.

Parameter	Value
$NN_{min}(NN_{max})$	0 (40)
N_{init}	5
$DN_{min}(DN_{max})$	1 (50)
ND_{init}	5
NDS_{pre}	4
NDS_{whi}	1
N_{ep}	10
MN_{inc}	0.03
I_u	-0.6
O_l	0.8
α	1.5
'Pre' development parameters	
$NH_{death}(NH_{birth})$	-0.405 (0.406)
$DH_{death}(DH_{birth})$	-0.39 (-0.197)
δ_{sh}	0.1
δ_{sp}	0.1487
δ_{sb}	0.07
δ_{dh}	0.1
δ_{dp}	0.1
δ_{dw}	0.101
'While' development parameters	
$NH_{death}(NH_{birth})$	-0.8 (0.8)
$DH_{death}(DH_{birth})$	-0.7 (0.7)
δ_{sh}	0.0011
δ_{sp}	0.0011
δ_{sb}	0.0011
δ_{dh}	0.0011
δ_{dp}	0.0011
δ_{dw}	0.0011

Twenty evolutionary runs were carried out using these parameters and the results are shown in Table 12. Once again, we see that parameter values are very precise. They were obtained using the same greedy search discussed earlier.

Table 12 Test accuracy for ten learning epochs versus no learning epochs.

Acc.	Learning epochs	No learning epochs
Mean	0.6974	0.7075
Median	0.6985	0.7266
Maximum	0.7355	0.7669
Minimum	0.6685	0.6200

Table 13 Test accuracies over problems using ten learning epochs .

Acc.	Cancer	Diabetes	Glass
Mean	0.8799	0.6250	0.3774
Median	0.9425	0.6354	0.3962
Maximum	0.9885	0.7031	0.5660
Minimum	0.6092	0.4740	0.0000

In Table 12 shows the results with multiple learning epochs versus those with no learning epochs. Table 13 shows the results on each problem using multiple learning epochs. It is clear that using no learning epochs gives better results. However, the results with multiple learning epochs are still reasonable despite the fact that the task is much more difficult, since one is effectively trying to evolve a *learning algorithm*. It is possible that further experimentation with developmental parameters could produce better results with multiple epochs.

Fig. 7 Variation of test classification accuracy with learning epoch for run 1 of 20

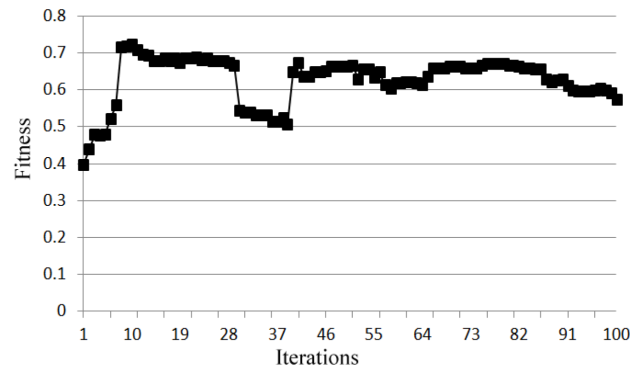

In Figure 7 we examine how the accuracy of the classifications varies with learning epochs (run 1 of 20). We set the maximum number of epochs to 100 now to see if learning continues beyond the upper limit used during evolution (10). We can see that classification accuracy increases with each epoch up to 10 epochs and starts to gradually decline after that. However, at epoch 29 the accuracy suddenly drops to 0.544 and at epoch 39 the accuracy increases again to 0.647. After this the accuracy shows a slow decline. We obtained several evolved solutions in which training

accuracy increased at each epoch until the imposed maximum number of epochs, however, as yet none of these were able to improve beyond the limit.

13 Open questions

There are many issues and questions that remain to be investigated.

Firstly, it is unclear why better results can not at present be obtained when evolving developmental programs with multiple epochs. Neither is it clear why programs can be evolved that continuously improve the developed ANNs over a number of epochs (i.e. 10) yet do not improve subsequently. It is worth contrasting the model discussed in this chapter with previous work on Self-Modifying CGP (SMCGP) [25]. In SMCGP phenotypes can be iterated to produce a sequence of programs or phenotypes. In some cases genotypes were found that produced general solutions and always improved at each iteration. The fitness was *accumulated* over all the correct test cases summed over all the iterations. In the problems studied (i.e. even-n parity, producing π) there was also a notion of perfection. For instance in the parity case perfection meant that at each iteration it produced the next parity case (with more inputs) perfectly. If at the next iteration, the appropriate parity function was not produced, then the iteration stopped. In the work discussed here, the fitness is not cumulative. At each epoch, the fitness is the average accuracy of the classifiers over the three classification problems. If the fitness reduces at the next epoch, then the epoch loop is terminated. However, in principle, we could sum the accuracies at each epoch and if an accuracy at a particular epoch is reduced, terminate the epoch loop. Summing the accuracies would give reward to developmental programs that produced the best *history* of developmental changes.

At present, the developmental programs do not receive a reward signal during multiple epochs. This means that the task for evolution is to continuously improve developed ANNs without being supplied with a reward signal. However, one would expect that as the fitness increases at each epoch the number of changes that need to be made to the developed ANNs should decrease. This suggests that supplying the fitness at the previous epoch to the developmental programs might be useful. In fact this option has already been implemented but as yet evidence is inconclusive that this produces improved results.

While learning over multiple epochs, we have assumed that the developmental parameters should be fixed (i.e. they are chosen before the development loop - see line 5 of Overview algorithm 1). However, it is not clear that this should be so. One could argue that during early learning topological changes in the brain network are more important and weight changes more important in later phases of learning. This suggests that at each step of the learning loop one could load developmental parameters, this would allow control of each epoch of learning. However, this has the drawback of increasing the number of possible parameter settings.

The neural variables that are given as inputs to the CGP developmental programs are an assumption of the model. For the soma these are: health, bias, position, prob-

lem type and average dendrite health, position and weight. For the dendrite they are: dendrite health, weight, position, problem type and soma health, bias and position. Further experimental work needs to be undertaken to determine whether they all are useful. The program written already has the inclusion of any of these variables as an option.

There are also many assumptions made in quite small aspects of the whole model. When new neurons or dendrites are born what should the initial values of the neural variables be? What are the best upper bounds on the number of neurons and dendrites? Currently, dendrite replication is decided by comparing the parent *neuron* health with DH_{birth} rather than comparing dendrite health with this threshold. If dendrite health was compared with a threshold it could rapidly lead to very large numbers of dendrites. Many choices have been made that need to be investigated in more detail. In real biological brains, early in the developmental process neurons move (and have no or few dendrites) and later neurons do not move and have many or at least a number of dendrites. This is understandable as moving when you have dendrites is difficult (if not impossible) as they provided resistance and would get obstructed by the dendrites of other neurons. However, in the proposed model movement can happen irrespective of such matters. It would be possible to restrict movement of whole neurons by making the movement increments depend on the number of dendrites a neuron has.

There are also very many parameters in the model and experiment has shown that results can be very sensitive to some of these. Thus further experimentation is required to identify good choices for these parameters.

A fundamental issue is how to handle inputs and outputs. In the classification problems the number of inputs is given by the problem with the most attributes, problems with less are given the value zero for those inputs. This could be awkward if the problems have hugely varying numbers of inputs. Is there another way of handling this? Perhaps one could borrow more ideas from SMCGP and make all input connections access inputs using pointer to a circular register of inputs. Every time a neuron connected to an input, a global pointer to the register of inputs would be incremented. Another possible idea is to assign all inputs a unique position (across all problems) and introduce the appropriate inputs at the ANN extraction stage, this would remove the need to assign zero to non-existent inputs (as mentioned above). This would mean no inputs are shared. Equally fundamental is the issue of handling outputs. Currently, we have dedicated output neurons for each output, however, this means that development can not start with a single neuron. Perhaps, neurons could decide to be an output neuron for a particular problem and some scheme would need to be devised to allocate the appropriate number of outputs for each computational problem (rather like was done in SMCGP). Alternatively, extra genes could be added to the genome like output genes in standard CGP. These output connection genes would be real-valued between -1 and 1 and snap to nearest neurons. Essentially, this would mean that the model would have three chromosomes, one each for the soma, dendrites and outputs. This would have the advantage that only non-output neurons would be necessary.

So far, we have examined the utility of the developmental model on three classification problems. However, the aim of the work is to produce general problem solving on many different kinds of computational problems. Clearly, a favourable direction to go is to expand the list of problems and problem types. How much neuron sharing would take place across problems of different types (e.g. classification and real-time control)? Would different kinds of problems cause whole new sub-networks to grow? These questions relate to a more fundamental issue which is the assessment of developmental ANNs. Should we have a training set of problems (rather than data) and evaluate on an unseen (but related) set of problems?

Currently the neurons exist in a one-dimensional space however it would be relatively straightforward to extend it to two or even three spatial dimensions.

In brains the morphology of neurons is activity dependent [53]. A simple way to introduce this would be to examine whether a neuron is actually involved in the propagation of signal from inputs to outputs (i.e. whether it is redundant or not). This activity could be input to developmental programs. Alternatively, a *signal related* input could be provided to developmental programs. This would mean that signals from other neurons in the model could influence decisions made by neural and dendrite programs. However, running developmental programs during the process of passing signals through the ANN would mean that conventional ANNs could not be extracted and also it would slow down assessment of network response to applied signals. Perhaps some statistical measures of signals could be computed which are supplied to neural programs. They could be calculated during each developmental step and then supplied to neuron and dendrite programs at the start of the next developmental step. However, it should be noted that activity dependent morphology implies that networks would change during training (and testing) and network morphology and behaviour would depend on past training history. This would complicate fitness assessment!

Eventually, the aim is to create developmental networks of spiking neurons. This would allow models of activity dependent development based on biological neurons to be abstracted and included in artificial models.

14 Conclusions

We have presented a conceptually simple model of a developmental neuron in which neural networks develop over time. Conventional ANNs can be extracted from these networks. We have shown that an evolved pair of programs can produce networks that can solve multiple classification problems reasonably well. Multiple-problem solving is a new domain for investigating more general developmental neural models.

15 Appendix: Detailed algorithms

15.1 Developing the brain and evaluating the fitness

The detailed algorithm for developing the brain and assessing its fitness is shown in Alg. 1 There are two stages to development. The first (which we refer to as 'pre') occurs prior to a learning epoch loop (lines 3-6). While the second phase (referred to as 'while') occurs inside a learning epoch loop (lines 9-12).

Lines 13-22 are concerned with calculating fitness. For each computational problem an ANN is extracted from the underlying brain. This is carried by a function $ExtractANN(problem, OutputAddress)$ which is detailed in Alg. 9. This function extracts a feedforward ANN corresponding to each computional problem (this is stored in a phenotype which we do not detail here). The array $OutputAddress$ stores the addresses of the output neurons associated with the computational problem. It is used together with the phenotype to extract the network of neurons that are required for the computational problem. Then the input data is supplied and the outputs of the ANN calculated. The class of a data instance is determined by the largest output. The learning loop (lines 8-29) develops the brain and exits if the fitness value (in this case classification accuracy) reduces (lines 23-27 in Alg. 1). One can think of the 'pre' development phase as growing a neural network prior to training. The 'while' phase is a period of development within the learning phase. N_{ep} denotes the user-defined number of learning epochs. N_p represents the number of problems in the suite of problems being solved. $N_{ex}(p)$ denotes the number of examples for each problem. A is the accuracy of prediction for a single training instance. F is the fitness over all examples. TF is the accumulated fitness over all problems. Fitness is normalised (lines 20 and 22).

Algorithm 1 Develop network and evaluate fitness

```
 1: function FITNESS
 2:     Initialise brain
 3:     Use 'pre' parameters
 4:     for s = 0 to s < NDS_pre do                # develop prior to learning
 5:         UpdateBrain
 6:     end for
 7:     TF_prev = 0
 8:     for e = 0 to e < N_ep do                   # learning loop
 9:         Use 'while' parameters                 # learning phase
10:         for s = 0 to s < NDS_whi do
11:             UpdateBrain
12:         end for
13:         TF = 0                                 # initialise total fit
14:         for p = 0 to p < N_p do
15:             ExtractANN(p, OutputAddress)       # Get ANN for problem p
16:             F = 0                              # initialise fit
17:             for t = 0 to t < N_ex(p) do
18:                 F = F + Acc                    # sum acc. over instances
19:             end for
20:             TF = TF + F/N_ex(p)                # sum normalised acc. over problems
21:         end for
22:         TF = TF/N_p                            # normalise total fitness
23:         if TF < TF_prev then                   # has fitness reduced?
24:             TF = TF_prev                       # return previous fitness
25:             Break                              # terminate learning loop
26:         else
27:             TF_prev = TF                       # update previous fitness
28:         end if
29:     end for
30:     return TF
31: end function
```

15.2 Developing the brain and evaluating the fitness

Algorithm 2 shows the update brain process. This algorithm is run at each developmental step. It runs the soma and dendrite programs for each neuron and from the previously existing brain creates a new version (*NewBrain*) which eventually overwrites the previous brain at the last step (lines 52-53).

Algorithm 2 Update brain

1: **function** UPDATEBRAIN
2: NewNumNeurons = 0
3: **for** $i = 0$ to $i <$ NumNeurons **do** # get number and addresses of neurons
4: **if** (Brain[i].out = 0) **then**
5: NonOutputNeuronAddress[NumNonOutputNeurons] = i
6: increment NumNonOutputNeurons
7: **else**
8: OutputNeuronAddress[NumOutputNeurons] = i
9: increment NumOutputNeurons
10: **end if**
11: **end for**
12: **for** $i = 0$ to $i <$ NumNonOutputNeurons **do** # process non-output neurons
13: NeuronAddress = NonOutputNeuronAddress[i]
14: Neuron = Brain[NeuronAddress]
15: UpdatedNeurVars = RunSoma(Neuron) # get new position, health and bias
16: **if** (DisallowNonOutputsToMove) **then**
17: UpdatedNeurVars.x = Neuron.x
18: **else**
19: UpdatedNeurVars.x = IfCollision(NewNumNeurons,NewBrain,UpdatedNeurVars.x)
20: **end if**
21: UpdatedNeuron = RunAllDendrites(Neuron, UpdatedNeurVars)
22: **if** (UpdatedNeuron.health $> NH_{death}$) **then** # if neuron survives
23: NewBrain[NewNumNeurons] = UpdatedNeuron
24: Increment NewNumNeurons
25: **if** (NewNumNeurons = NN_{max}-NumOutputNeurons) **then**
26: **Break** # exit non-output neuron loop
27: **end if**
28: **end if**
29: **if** (UpdatedNeuron.health $> NH_{health}$) **then** # neuron replicates
30: UpdatedNeuron.x = UpdatedNeuron.x$+MN_{inc}$
31: UpdatedNeuron.x = IfCollision(NewNumNeurons, NewBrain, UpdatedNeuron.x)
32: NewBrain[NewNumNeurons] = CreateNewNeuron(UpdatedNeuron)
33: Increment NewNumNeurons
34: **if** (NewNumNeurons = NN_{max} - NumOutputNeurons) **then**
35: **Break** # exit non-output neuron loop
36: **end if**
37: **end if**
38: **end for**
39: **for** $i = 0$ to $i <$ NumOutputNeurons **do** # process output neurons
40: NeuronAddress = OutputNeuronAddress[i]
41: Neuron = Brain[NeuronAddress]
42: UpdatedNeurVars = RunSoma(Neuron) # get new position, health and bias
43: **if** (DisallowOutputsToMove) **then**
44: UpdatedNeurVars.x = Neuron.x
45: **else**
46: UpdatedNeurVars.x = IfCollision(NewNumNeurons,NewBrain,UpdatedNeurVars.x)
47: **end if**
48: UpdatedNeuron = RunAllDendrites(UpdatedNeuron)
49: NewBrain[NewNumNeurons] = UpdatedNeuron
50: Increment NewNumNeurons
51: **end for**
52: NumNeurons = NewNumNeurons
53: Brain = NewBrain
54: **end function**

Alg. 2 starts by analyzing the brain to determine the addresses and numbers of non-output and output neurons (lines 3-11). Then the non-output neurons are processed. The evolved soma program is executed and it returns a neuron with updated values for the neuron position, health and bias. These are stored in the variable $UpdatedNeurVars$.

If the user-defined option to disallow non-output neuron movement is chosen then the updated neuron position is reset to that before the soma program is run (lines 16-17). Next the evolved dendrite programs are executed in all dendrites. The algorithmic details are given in Alg. 6 (See Sect. 4.6).

The neuron health is compared with the user-defined neuron death threshold NH_{death} and if the health exceeds the threshold the neuron survives (see lines 22-28). At this stage it is possible that the neuron has been given a position that is identical to one of the neurons in the developing brain (*NewBrain*) so one needs a mechanism for preventing this. This is accomplished by Alg. 3 (Lines 19 and 46). It checks whether a collision has occurred and if so an increment MN_{inc} is added to the position and then it is bound to the interval [-1, 1]. In line 23 the updated neuron is written into *NewBrain*. A check is made in line 25 to see if the allowed number of neurons has been reached, if so the non output neuron update loop (lines 12 to 38) is exited and the output neuron section starts (lines 39 to 51). If the limit on numbers of neurons has not been reached, the updated neuron may replicate depending on whether its health is above the user-defined threshold, NH_{health} (line 29). The position of the new born neuron is immediately incremented by MN_{inc} so that it does not collide with its parent (line 30). However, its position needs to be checked also to see if it collides with any other neuron, in which case its position is incremented again until a position is found that causes no collision. This is done in the function IFCOLLISION.

In CREATENEWNEURON (see line 32) the bias, the incremented position and dendrites of the parent neuron are copied into the child neuron. However, the new neuron is given a health of 1.0 (the maximum value). The algorithm examines the non-output neurons (lines 39-51) and again is terminated if the allowed number of neurons is exceeded. The steps are similar to those carried out with non-output neurons, except that output neurons can not either die or replicate as their number is fixed by the number of outputs required by the computational problem being solved.

The details of the neuron collision avoidance mechanism is shown in Alg. 3.

15.3 Running the soma

The UPDATEBRAIN program calls the RUNSOMA program (Alg. 4) to determine how the soma changes in each developmental step. The seven soma program inputs comprising the neuron health, position and bias, the averaged position, weight and health of the neuron's dendrites and the problem type are supplied to the CGP encoded soma program (line 12). The array *ProblemTypeInputs* stores

Algorithm 3 Move neuron if it collides with another.

```
 1: function IFCOLLISION(NumNeurons, Brain, NeuronPosition)
 2:     NewPosition = NeuronPosition
 3:     collision = 1
 4:     while collision do
 5:         collision = 0
 6:         for i = 0 to j < NumNeurons do
 7:             if (| NeuronPosition - Brain[i].x | < 1.e-8) then
 8:                 collision = 1
 9:             end if
10:             if collision then
11:                 break
12:             end if
13:         end for
14:         if collision then
15:             NewPosition = NewPosition+MN_inc
16:         end if
17:     end while
18:     if collision then
19:         NewPosition = Bound(NewPosition)
20:     end if
21:     return NewPosition
22: end function
```

NumProblems+1 constants equally spaced between -1 and 1. These are used to allow output neurons to know what computational problem they belong to.

The soma program has three outputs relating to the position, health and bias of the neuron. These are used to update the neuron (line 13).

Algorithm 4 RunSoma(Neuron)

```
 1: function RUNSOMA(Neuron)
 2:     AvDendritePosition = GetAvDendritePosition(Neuron)
 3:     AvDendriteWeight = GetAvDendriteWeight(Neuron)
 4:     AvDendriteHealth = GetAvDendriteHealth(Neuron)
 5:     SomaProgramInputs[0] = Neuron.health
 6:     SomaProgramInputs[1] = Neuron.x
 7:     SomaProgramInputs[2] = Neuron.bias
 8:     SomaProgramInputs[3] = AvDendritePosition
 9:     SomaProgramInputs[4] = AvDendriteWeight
10:     SomaProgramInputs[5] = AvDendriteHealth
11:     SomaProgramInputs[6] = ProblemTypeInputs[WhichProblem]
12:     SomaProgramOutputs = SomaProgram(SomaProgramInputs)
13:     UpdatedNeuron = UpdateNeuron(Neuron, SomaProgramOutputs)
14:     return UpdatedNeuron.x, UpdatedNeuron.health, UpdatedNeuron.bias
15: end function
```

15.4 Changing the Neuron Variables

The UPDATENEURON algorithm (5) updates the neuron properties of health, position and bias according to three user-chosen options defined by a variable Incr_{opt}. If this is zero, then the soma program outputs determine directly the updated values of the soma's health, position and bias. If Incr_{opt} is one or two, the updated values of the soma are changed from the parent neuron's values in an incremental way. This is either a linear or nonlinear increment or decrement depending on whether the soma program's outputs are greater than or less than or equal to zero (lines 8 to 16). The magnitudes of the increments is defined by the user-defined constants: δ_{sh}, δ_{sp}, δ_{sb} and sigmoid slope parameter, α (see Table 1).

The increment methods described in Algorithm 5 change neural variables, so action needs to be taken to force the variables to strictly lie in the interval $[-1, 1]$. We call this 'bounding' (lines 34-36).This is accomplished using a hyperbolic tangent function.

15.5 Running all dendrite programs and building a new neuron

Alg. 6 takes an existing neuron and creates a new neuron using the updated soma variables, position, health and bias which are stored in $UpdateNeurVars$ (from Alg. 4) and the updated dendrites which result from running the dendrite program in all the dendrites. Initially (line 3-5), the updated soma variables are written into the updated neuron. The number of dendrites in the updated neuron is set to zero. In lines 8-11, the health of the non-updated *neuron* is examined and if it is above the *dendrite health threshold* for birth, a new dendrite is generated and the updated neuron gains a dendrite. If so, the neuron gains a dendrite created by a function $GenerateDendrite()$. This assigns a weight, health and position to the new dendrite. The weight and health is set to one and the position set to half the parent neuron position. These choices appeared to give good results.

Lines 12-33 are concerned with processing the dendrite program in all the dendrites of the non-updated neuron and updating the dendrites. If the updated dendrite has a health above its death threshold then it survives and gets written into the updated neuron (lines 22-28). Updated dendrites do not get written into the updated neuron if it already has the maximum allowed number of dendrites (line 25-27). In lines 30-33 a check is made as to whether the updated neuron has no dendrites. If this is so, it is given one of the dendrites of the non-updated neuron. Finally, the updated neuron is returned to the calling function.

Alg. 6 calls the function RUNDENDRITE (line 21). This function is detailed in Alg. 7. It changes the dendrites of a neuron according to the evolved dendrite program. It begins by assigning the dendrites health, position and weight to the parent dendrite variables. It writes the dendrite program outputs to the internal variables health, weight and position. Then in lines 8-16 it defines the possible increments in health, weight and position that will be used to increment or decrement the parent

Algorithm 5 Neuron update function

```
 1: function UPDATENEURON(Neuron, SomaProgramOutputs)
 2:     ParentHealth = Neuron.health
 3:     ParentPosition = Neuron.x
 4:     ParentBias = Neuron.bias
 5:     health = SomaProgramOutputs[0]
 6:     position = SomaProgramOutputs[1]
 7:     bias = SomaProgramOutputs[2]
 8:     if (Incr_opt = 1) then                       # calculate increment
 9:         HealthIncrement = δ_sh
10:         PositionIncrement = δ_sp
11:         BiasIncrement = δ_sb
12:     else if (Incr_opt = 2) then
13:         HealthIncrement = δ_sh*sigmoid(health, α)
14:         PositionIncrement = δ_sp*sigmoid(position, α)
15:         BiasIncrement = δ_sb*sigmoid(bias, α)
16:     end if
17:     if (Incr_opt > 0) then                       # apply increment
18:         if (health > 0.0) then
19:             health = ParentHealth + HealthIncrement
20:         else
21:             health = ParentHealth - HealthIncrement
22:         end if
23:         if (position > 0.0) then
24:             position = ParentPosition + PositionIncrement
25:         else
26:             health = ParentPosition - PositionIncrement
27:         end if
28:         if (bias > 0.0) then
29:             bias = ParentBias + BiasIncrement
30:         else
31:             bias = ParentBias - BiasIncrement
32:         end if
33:     end if
34:     health = Bound(health)
35:     position = Bound(position)
36:     bias = Bound(bias)
37:     return health, position and bias
38: end function
```

variables according to the user defined incremental options (linear or non-linear).In lines 17-33 it respectively carries out the increments or decrements of the parent dendrite variables according whether the corresponding dendrite program outputs are greater than or less than or equal to zero. After this it bounds those variables. Finally, in lines 37-44 it updates the dendrites health, weight and position provided the adjusted health is above the dendrite death threshold (in other words it survives). Note that if $Incr_{opt} = 0$ then there is no incremental adjustment and the health, weight and position of the dendrites are just bounded (lines 34-36).

Algorithm 6 Run the evolved dendrite program in all dendrites

```
1: function RUNALLDENDRITES(Neuron, DendriteProgram, NewSomaPosition, NewSoma-
   Health, NewSomaBias)
2:     WhichProblem = Neuron.isout
3:     OutNeuron.x = NewSomaPosition
4:     OutNeuron.health = NewSomaHealth
5:     OutNeuron.bias = NewSomaBias
6:     OutNeuron.isout = WhichProblem
7:     OutNeuron.NumDendrites = 0
8:     if (Neuron.health > DH_birth ) then
9:         OutNeuron.dendrites[NumDendrites] = GenerateDendrite()
10:        Increment OutNeuron.NumDendrites
11:    end if
12:    for i = 0 to i < OutNeuron.NumDendrites do
13:        DendriteProgramInputs[0] = Neuron.health
14:        DendriteProgramInputs[1] = Neuron.x
15:        DendriteProgramInputs[2] = Neuron.bias
16:        DendriteProgramInputs[3] = Neuron.dendrites[i].health
17:        DendriteProgramInputs[4] = Neuron.dendrites[i].weight
18:        DendriteProgramInputs[5] = Neuron.dendrites[i].position
19:        DendriteProgramInputs[6] = ProblemTypeInputs[WhichProblem]
20:        DendriteProgramOutputs = DendriteProgram(DendriteProgramInputs)
21:        UpdatedDendrite = RunDendrite(Neuron, DendriteProgramOutputs)
22:        if (UpdatedDendrite.isAlive) then
23:            OutNeuron.dendrites[NumDendrites] = UpdatedDendrite
24:            increment OutNeuron.NumDendrites
25:            if (OutNeuron.NumDendrites > MaxNumDendrites) then
26:                break
27:            end if
28:        end if
29:    end for
30:    if (OutNeuron.NumDendrites = 0) then          # if all dendrites die
31:        OutNeuron.dendrites[0] = Neuron.dendrites[0]
32:        OutNeuron.NumDendrites = 1
33:    end if
34:    return OutNeuron
35: end function
```

Alg. 2 uses a function CREATENEWNEURON to create a new neuron if the neuron health is above a threshold. This function is described in Alg. 8. It makes the new born neuron the same as the parent (note, its position will be adjusted by the collision avoidance algorithm) except that it is given a health of one. Experiments suggested that this gave better results.

Algorithm 7 Change dendrites according to the evolved dendrite program

```
1:  function RunDendrite(Neuron, WhichDendrite, DendriteProgramOutputs)
2:      ParentHealth = Neuron.dendrites[WhichDendrite].health
3:      ParentPosition = Neuron.dendrites[WhichDendrite].x
4:      ParentWeight = Neuron.dendrites[WhichDendrite].weight
5:      health = DendriteProgramOutputs[0]
6:      weight = DendriteProgramOutputs[1]
7:      position = DendriteProgramOutputs[2]
8:      if (Incr_opt = 1) then
9:          HealthIncrement = δ_dh
10:         WeightIncrement = δ_dw
11:         PositionIncrement = δ_dp
12:     else if (Incr_opt = 2) then
13:         HealthIncrement = δ_dh*sigmoid(health, α)
14:         WeightIncrement = δ_dw*sigmoid(weight, α)
15:         PositionIncrement = δ_dp*sigmoid(position, α)
16:     end if
17:     if (Incr_opt > 0) then
18:         if (health > 0.0) then
19:             health = ParentHealth + HealthIncrement
20:         else
21:             health = ParentHealth - HealthIncrement
22:         end if
23:         if (position > 0.0) then
24:             position = ParentPosition + PositionIncrement
25:         else
26:             health = ParentPosition - PositionIncrement
27:         end if
28:         if (weight > 0.0) then
29:             weight = ParentWeight + BiasIncrement
30:         else
31:             weight = ParentWeight - BiasIncrement
32:         end if
33:     end if
34:     health = Bound(health)
35:     position = Bound(position)
36:     weight = Bound(weight)
37:     if (health > DH_death) then
38:         UpdatedDendrite.weight = weight
39:         UpdatedDendrite.health = health
40:         UpdatedDendrite.x = position
41:         UpdatedDendriteisAlive = 1
42:     else
43:         UpdatedDendriteisAlive = 0
44:     end if
45:     return UpdatedDendrite and UpdatedDendriteisAlive
46: end function
```

15.6 Extracting conventional ANNs from the evolved brain

In algorithm 1, a conventional feed-forward ANN is extracted from the underlying collection of neurons (line 15). The algorithm for doing this is shown in algorithm 9.

Algorithm 8 Create new neuron from parent neuron

```
 1: function CREATENEWNEURON(ParentNeuron)
 2:     ChildNeuron.NumDendrites = ParentNeuron.NumDendrites
 3:     ChildNeuron.isout = 0
 4:     ChildNeuron.health = 1
 5:     ChildNeuron.bias = ParentNeuron.bias
 6:     ChildNeuron.x = ParentNeuron.x
 7:     for i = 0 to i < ChildNeuron.NumDendrites do
 8:         ChildNeuron.dendrites[i] = ParentNeuron.dendrites[i]
 9:     end for
10: end function
```

Firstly, this algorithm determines the number of inputs to the ANN (line 5). Since inputs are shared across problems the number of inputs is set to be the maximum number of inputs that occur in the computational problem suite. If an individual problem has less inputs than this maximum, the extra inputs are set to 0.0. The brain array is sorted by position. The algorithm then examines all neurons (line 7) and calculates the number of non-output neurons and output neurons and stores the neuron data in arrays *NonOutputNeurons* and *OutputNeurons*. It also calculates their addresses in the brain array.

The next phase is to go through all dendrites of the non-output neurons to determine which inputs or neurons they connect to (lines 19 to 33). The evolved neuron programs generate dendrites with end positions anywhere in the interval [-1, 1]. The end positions are converted to lengths (line 25). In this step the dendrite position is linearly mapped into the interval [0, 1]. To generate a valid neural network we assume that dendrites are automatically connected to the nearest neuron or input *on the left*. We refer to this as "snapping" (lines 28 and 44). The dendrites of non-output neurons are allowed to connect to either inputs or other non-output neurons on their left. However, output neurons are only allowed to connect to non-output neurons on their left. Algorithm 10 returns the address of the neuron or input that the dendrite snaps to. The dendrites of output neurons are not allowed to connect directly to inputs (see Line 4 of the GETCLOSEST function), however, when neurons are allowed to move, there can occur a situation where an output neuron is positioned so that it is the first neuron on the right of the outputs. In that situation it can only connect to inputs. If this situation occurs then the initialisation of the variable *AddressOfClosest* to zero in the GETCLOSEST function (line 2) means that all the dendrites of the output neuron will be connected to the first external input to the ANN network. Thus a valid network will still be extracted albeit with a rather useless output neuron. It is expected that evolution will avoid using programs that allow this to happen.

Algorithm 9 stores the information required to extract the ANN in an array called *Phenotype*. It contains the connection addresses of all neurons and their weights (lines 29-30 and 45-46). Finally it stores the addresses of the output neurons (*OutputAddress*) corresponding to the computational problem whose ANN is being extracted (lines 49-52). These define the outputs of the extracted ANNs when

they are supplied with inputs (i.e. in the fitness function when the Accuracy is assessed (see Alg. 1). The *Phenotype* is stored in the same format as Cartesian Genetic Programming (see section 5) and decoded in a similar way to genotypes.

References

1. Aljundi, R., Chakravarty, P., Tuytelaars, T.: Expert gate: Lifelong learning with a network of experts. CoRR, abs/1611.06194 **2** (2016)
2. Astor, J.C., Adami, C.: A development model for the evolution of artificial neural networks. Artificial Life **6**, 189–218 (2000)
3. Balaam, A.: Developmental neural networks for agents. In: Advances in Artificial Life, Proceedings of the 7th European Conference on Artificial Life (ECAL 2003), pp. 154–163. Springer (2003)
4. Belew, R.K.: Interposing an ontogenic model between genetic algorithms and neural networks. In: S.J. Hanson, J.D. Cowan, C.L. Giles (eds.) Advances in neural information processing systems NIPS5, pp. 99–106. Morgan Kaufmann (1993)
5. Boers, E.J.W., Kuiper, H.: Biological metaphors and the design of modular neural networks. Master's thesis, Dept. of Comp. Sci. and Dept. of Exp. and Theor. Psych., Leiden University (1992)
6. Cangelosi, A., Nolfi, S., Parisi, D.: Cell division and migration in a 'genotype' for neural networks. Network-Computation in Neural Systems **5**, 497–515 (1994)
7. Deacon, T.: The Symbolic Species: The Co-evolution of Language and the Brain. W.W. Norton and Company, New York (1998)
8. Dekaban, A.S., Sadowsky, D.: Changes in brain weights during the span of human life. Ann. Neurol. **4**, 345–356 (1978)
9. Downing, K.L.: Supplementing evolutionary developmental systems with abstract models of neurogenesis. In: Proc. Conf. on Genetic and evolutionary Comp., pp. 990–996 (2007)
10. Drchal, J., Šnorek, M.: Tree-based indirect encodings for evolutionary development of neural networks. In: V. Kůrková, R. Neruda, J. Koutník (eds.) Artificial Neural Networks - ICANN, pp. 839–848. Springer Berlin Heidelberg, Berlin, Heidelberg (2008)
11. Edelman, G., Tononi, G.: A Universe of Consciousness. Basic Books, New York (2000)
12. Eggenberger, P.: Creation of neural networks based on developmental and evolutionary principles. In: W. Gerstner, A. Germond, M. Hasler, J.D. Nicoud (eds.) Artificial Neural Networks — ICANN'97, pp. 337–342 (1997)
13. Fahlman, S.E., Lebiere, C.: The cascade-correlation learning architecture. In: Advances in neural information processing systems, pp. 524–532 (1990)
14. Federici, D.: A regenerating spiking neural network. Neural Networks **18**(5-6), 746–754 (2005)
15. Fernández-Delgado, M., Cernadas, E., Barro, S., Amorim, D.: Do we need hundreds of classifiers to solve real world classification problems? J. Mach. Learn. Res. **15**(1), 3133–3181 (2014)
16. Ferreira, C.: Gene Expression Programming: Mathematical Modeling by an Artificial Intelligence, 2 edn. Springer, New York (2006)
17. Floreano, D., Urzelai, J.: Neural morphogenesis, synaptic plasticity, and evolution. Theory in Biosciences **120**(3), 225–240 (2001)
18. Franco, L., Jerez, J.M.: Constructive neural networks, vol. 258. Springer (2009)
19. French, R.M.: Catastrophic Forgetting in Connectionist Networks: Causes, Consequences and Solutions. Trends in Cognitive Sciences **3**(4), 128–135 (1999)
20. Goldman, B.W., Punch, W.F.: Reducing wasted evaluations in cartesian genetic programming. In: Genetic Programming: 16th European Conference, EuroGP 2013, Vienna, Austria, April 3-5, 2013. Proceedings, pp. 61–72. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)

21. Goldman, B.W., Punch, W.F.: Analysis of cartesian genetic programmings evolutionary mechanisms. Evolutionary Computation, IEEE Transactions on **19**, 359 – 373 (2015)
22. Gruau, F.: Automatic definition of modular neural networks. Adaptive Behaviour **3**, 151–183 (1994)
23. Gruau, F., Whitley, D., Pyeatt, L.: A comparison between cellular encoding and direct encoding for genetic neural networks. In: Proc. Conf. on Genetic Programming, pp. 81–89 (1996)
24. Hampton, A.N., Adami, C.: Evolution of robust developmental neural networks. In: J. Pollack, M.A. Bedau, P. Husbands, T. Ikegami, R. Watson (eds.) Proceedings of Artificial Life IX, pp. 438–443 (2004)
25. Harding, S., Miller, J.F., Banzhaf, W.: Developments in cartesian genetic programming: Self-modifying cgp. Genetic Programming and Evolvable Machines **11**(3-4), 397–439 (2010)
26. Hornby, G., Lipson, H., Pollack, J.B.: Generative representations for the automated design of modular physical robots. IEEE Trans. on Robotics and Automation **19**, 703–719 (2003)
27. Hornby, G.S., Pollack, J.B.: Creating high-level components with a generative representation for body-brain evolution. Artificial Life **8(3)** (2002)
28. Huizinga, J., Clune, J., Mouret, J.B.: Evolving neural networks that are both modular and regular: HyperNEAT plus the connection cost technique. In: Proc. Conf. on Genetic and Evolutionary Computation, pp. 697–704 (2014)
29. Isles, A.: Neural and Behavioral Epigenetics; what it Is, and what is Hype. John Wiley & Sons Limited (2015)
30. Jakobi, N.: Harnessing Morphogenesis, COGS Research Paper 423. Tech. rep., University of Sussex (1995)
31. Jung, S.Y.: A topographical method for the development of neural networks for artificial brain evolution. Artificial Life **11**, 293–316 (2005)
32. Kandel, E.R., Schwartz, J.H., Jessell: Principles of Neural Science, 4th Edition. McGraw-Hill (2000)
33. Khan, G.M.: Evolution of Artificial Neural Development - In Search of Learning Genes, *Studies in Computational Intelligence*, vol. 725. Springer (2018)
34. Khan, G.M., Miller, J.F.: In search of intelligence: evolving a developmental neuron capable of learning. Connect. Sci. **26**(4), 297–333 (2014)
35. Khan, G.M., Miller, J.F., Halliday, D.M.: Evolution of Cartesian Genetic Programs for Development of Learning Neural Architecture. Evol. Computation **19**(3), 469–523 (2011)
36. Kitano, H.: Designing neural networks using genetic algorithms with graph generation system. Complex Systems **4**, 461–476 (1990)
37. Kleim, J.A., Lussnig, E., Schwartz, E.R., Comery, T.A., Greenough, W.T.: Synaptogenesis and fos expression in the motor cortex of the adult rat after motor skill learning. J. Neurosci **16**, 4529–4535 (1996)
38. Kleim, J.A., Vij, K., Ballard, D.H., Greenough, W.T.: Learning-dependent synaptic modifications in the cerebellar cortex of the adult rat persist for at least four weeks. J. Neurosci **17**, 717–721 (1997)
39. Kodjabachian, J., Meyer, J.A.: Evolution and development of neural controllers for locomotion, gradient-following, and obstacle-avoidance in artificial insects. IEEE Transactions on Neural Networks **9**, 796–812 (1998)
40. Koutník, J., Gomez, F., Schmidhuber, J.: Evolving neural networks in compressed weight space. In: Proc. Conference on Genetic and Evolutionary Computation (GECCO-10) (2010)
41. Kumar, S., Bentley, P. (eds.): On Growth, Form and Computers. Academic Press (2003)
42. Luke, S., Spector, L.: Evolving graphs and networks with edge encoding: Preliminary report. In: Late Breaking Papers at the Genetic Programming Conference, pp. 117–124 (1996)
43. Maguire, E.A., Gadian, D.G., Johnsrude, I.S., Good, C.D., Ashburner, J., Frackowiak, R.S.J., Frith, C.D.: Navigation-related structural change in the hippocampi of taxi drivers. PNAS **97**, 4398–4403 (2000)
44. McCloskey, M., Cohen, N.: Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem. The Psychology of Learning and Motivation **24**, 109–165 (1989)
45. McCulloch, W., Pitts, W.: A logical calculus of the ideas immanent in nervous activity. The Bulletin of Mathematical Biophysics **5**, 115–133 (1943)

46. Métin, C., Vallee, R., Rakic, P., Bhide, P.: Modes and mishaps of neuronal migration in the mammalian brain. Neuroscience **28**, 11,746–11,752 (2008)
47. Miller, J.F.: What bloat? cartesian genetic programming on boolean problems. In: Proc. Conf. Genetic and Evolutionary Computation, Late breaking papers, pp. 295–302 (2001)
48. Miller, J.F. (ed.): Cartesian Genetic Programming. Springer (2011)
49. Miller, J.F., Smith, S.L.: Redundancy and computational efficiency in Cartesian Genetic Programming. IEEE Trans. on Evolutionary Computation **10**(2), 167–174 (2006)
50. Miller, J.F., Thomson, P.: Cartesian genetic programming. In: Proc. European Conf. on Genetic Programming, *LNCS*, vol. 10802, pp. 121–132 (2000)
51. Miller, J.F., Thomson, P.: A Developmental Method for Growing Graphs and Circuits. In: Proc. Int. Conf. on Evolvable Systems, *LNCS*, vol. 2606, pp. 93–104 (2003)
52. Miller, J.F., Wilson, D.G., Cussat-Blanc, S.: Evolving developmental programs that build neural networks for solving multiple problems. In: W. Banzhaf, L. Spector, L. Sheneman (eds.) Genetic Programming Theory and Practice XVI, chap. TBC, p. TBC. Springer (2019)
53. Ooyen, A.V. (ed.): Modeling Neural Development. MIT Press (2003)
54. Rakic, P.: Principles of neural cell migration. Experientia **46**, 882–891 (1990)
55. Ratcliff, R.: Connectionist Models of Recognition and Memory: Constraints Imposed by Learning and Forgetting Functions. Psychological Review **97**, 205–308 (1990)
56. Risi, S., Lehman, J., Stanley, K.O.: Evolving the placement and density of neurons in the HyperNEAT substrate. In: Proc. Conf. on Genetic and Evolutionary Computation, pp. 563–570 (2010)
57. Risi, S., Stanley, K.O.: Indirectly encoding neural plasticity as a pattern of local rules. In: From Animals to Animats 11: Conf. on Simulation of Adaptive Behavior (2010)
58. Risi, S., Stanley, K.O.: Enhancing ES-HyperNEAT to evolve more complex regular neural networks. In: Proc. Conf. on Genetic and Evolutionary Computation, pp. 1539–1546 (2011)
59. Roggen, D., Federici, D., Floreano, D.: Evolutionary morphogenesis for multi-cellular systems. Genetic Programming and Evolvable Machines **8**(1), 61–96 (2007)
60. Rose, S.: The Making of Memory: From Molecules to Mind. Vintage (2003)
61. Rust, A., Adams, R., Bolouri, H.: Evolutionary neural topiary: Growing and sculpting artificial neurons to order. In: Proc. Conf. on the Simulation and synthesis of Living Systems, pp. 146–150 (2000)
62. Rusu, A.A., Rabinowitz, N.C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., Hadsell, R.: Progressive neural networks. arXiv preprint arXiv:1606.04671 (2016)
63. Ryan, C., Collins, J.J., Neill, M.O.: Grammatical evolution: Evolving programs for an arbitrary language. In: W. Banzhaf, R. Poli, M. Schoenauer, T.C. Fogarty (eds.) Genetic Programming, pp. 83–96. Springer Berlin Heidelberg (1998)
64. Sharkey, A.J.: Combining artificial neural nets: ensemble and modular multi-net systems. Springer Science & Business Media (2012)
65. Siddiqi, A.A., Lucas, S.M.: A comparison of matrix rewriting versus direct encoding for evolving neural networks. In: Proceedings IEEE International Conference on Evolutionary Computation Proceedings, pp. 392–397 (1998)
66. Smythies, J.: The Dynamic Neuron. MIT Press (2002)
67. Stanley, K., Miikkulainen, R.: Efficient evolution of neural network topologies. In: Proc. Congress on Evolutionary Computation, vol. 2, pp. 1757–1762 (2002)
68. Stanley, K.O.: Compositional pattern producing networks: A novel abstraction of development. Genetic Programming and Evolvable Machines **8**, 131–162 (2007)
69. Stanley, K.O., D'Ambrosio, D.B., Gauci, J.: A hypercube-based encoding for evolving large-scale neural networks. Artificial Life **15**, 185–212 (2009)
70. Stanley, K.O., Miikkulainen, R.: A taxonomy for artificial embryogeny. Artificial Life **9(2)**, 93–130 (2003)
71. Suchorzewski, M., Clune, J.: A novel generative encoding for evolving modular, regular and scalable networks. In: Proc. Conf. on Genetic and Evolutionary Computation, pp. 1523–1530 (2011)

72. Terekhov, A.V., Montone, G., ORegan, J.K.: Knowledge transfer in deep block-modular neural networks. In: Conference on Biomimetic and Biohybrid Systems, pp. 268–279. Springer (2015)
73. Tierney, A., Nelson III, C.: Brain development and the role of experience in the early years. Zero Three **30**, 9–13 (2009)
74. Tramontin, A.D., Brenowitz, E.: Seasonal plasticity in the adult brain. Trends in Neuroscience **23**, 251–258 (2000)
75. Tsankova, N., Renthal, W., Kumar, A., Nestler, E.: Epigenetic regulation in psychiatric disorders. Nature Reviews Neuroscience **8**(5), 33–367 (2007)
76. Turner, A.J., Miller, J.F.: Cartesian Genetic Programming encoded artificial neural networks: A comparison using three benchmarks. In: Proceedings of the Conference on Genetic and Evolutionary Computation (GECCO), pp. 1005–1012 (2013)
77. Turner, A.J., Miller, J.F.: Recurrent cartesian genetic programming. In: Proc. Parallel Problem Solving from Nature, pp. 476–486 (2014)
78. Valverde, F.: Rate and extent of recovery from dark rearing in the visual cortex of the mouse. Brain Res. **33**, 1–11 (1971)
79. Vassilev, V.K., Miller, J.F.: The Advantages of Landscape Neutrality in Digital Circuit Evolution. In: Proc. Int. Conf. on Evolvable Systems, *LNCS*, vol. 1801, pp. 252–263. Springer Verlag (2000)
80. Yerushalmi, U., Teicher, M.: Evolving synaptic plasticity with an evolutionary cellular development model. PLOS One **3**(11), e3697 (2008)
81. Yu, T., Miller, J.F.: Neutrality and the Evolvability of Boolean function landscape. In: Proc. European Conference on Genetic Programming, *LNCS*, vol. 2038, pp. 204–217 (2001)
82. Zar, J.H.: Biostatistical Analysis, 2nd edn. Prentice Hall (1984)

Algorithm 9 The extraction of neural networks from the underlying brain.

```
 1: function EXTRACTANN(problem, OutputAddress)
 2:     NumNonOutputNeurons = 0
 3:     NumOutputNeurons = 0
 4:     OutputCount=0
 5:     N_i = max(N_i, p)
 6:     sort(Brain, 0, NumNeurons-1)                # sort neurons by position
 7:     for i = 0 to i < NumNeurons do
 8:         Address = i + N_i
 9:         if (Brain[i].isout > 0) then            # non-output neuron
10:             NonOutputNeur[NumNonOutputNeur] = Brain[i]
11:             NonOutputNeuronAddress[NumNonOutputNeur]= Address
12:             Increment NumNonOutputNeur
13:         else                                    # output neuron
14:             OutputNeurons[NumOutputNeurons]= Brain[i]
15:             OutputNeuronAddress[NumOutputNeurons]= Address
16:             Increment NumOutputNeurons
17:         end if
18:     end for
19:     for i = 0 to i < NumNonOutputNeur do       # do non-output neurons
20:         Phenotype[i].isout = 0
21:         Phenotype[i].bias = NonOutputNeur[i].bias
22:         Phenotype[i].address = NonOutputNeuronAddress[i]
23:         NeuronPosition = NonOutputNeur[i].x
24:         for j = 0 to j < NonOutputNeur[i].NumDendrites do
25:             Convert DendritePosition to DendriteLength
26:             DendPos = NeuronPosition - DendriteLength
27:             DendPos = Bound(DendPos)
28:             AddressClosest = GetClosest(NumNonOutputNeur, NonOutputNeur, 0, DendPos)
29:             Phenotype[i].ConnectionAddresses[j] = AddressClosest
30:             Phenotype[i].weights[j] = NonOutputNeur[i].weight[j]
31:         end for
32:         Phenotype[i].NumConnectionAddress = NonOutputNeur[i].NumDendrites
33:     end for
34:     for i = 0 to i < NumOutputNeurons do       # do output neurons
35:         i1 = i+NumOutputNeurons
36:         Phenotype[i1].isout = OutputNeurons[i].isout
37:         Phenotype[i1].bias = OutputNeurons[i].bias
38:         Phenotype[i1].address = OutputNeuronAddress[i]
39:         NeuronPosition = OutputNeurons[i].x
40:         for j = 0 to j < OutputNeurons[i].NumDendrites do
41:             Convert DendritePosition to DendriteLength
42:             DendPos = NeuronPosition - DendriteLength
43:             DendPos = Bound(DendPos)
44:             AddressClosest = GetClosest(NumNonOutputNeur, NonOutputNeur, 1, DendPos)
45:             Phenotype[i1].ConnectionAddresses[j] = AddressClosest
46:             Phenotype[i1].weights[j] = OutputNeuron[i].weight[j]
47:         end for
48:         Phenotype[i1].NumConnectionAddress = OutputNeurons[i].NumDendrites
49:         if (OutputNeurons[i].isout == problem+1) then
50:             OutputAddress[OutputCount] = OutputNeuronAddress[i]
51:             Increment OutputCount
52:         end if
53:     end for
54: end function
```

Algorithm 10 Find which input or neuron a dendrite is closest to

```
 1: function GETCLOSEST(NumNonOutNeur, NonOutNeur, IsOut, DendPos)
 2:     AddressOfClosest = 0
 3:     min = 3.0
 4:     if (IsOut = 0) then                          # only non-out neurons connect to inputs
 5:         for (i = 0 to i < MaxNumInputs) do
 6:             distance = DendPos - InputLocations[i]
 7:             if distance > 0 then
 8:                 if (distance < min) then
 9:                     min = distance
10:                     AddressOfClosest = i
11:                 end if
12:             end if
13:         end for
14:     end if
15:     for j = 0 to j <NumNonOutputNeur do
16:         distance = DendPos - NonOutNeur[j].x
17:         if distance > 0 then                     # feed-forward connections
18:             if (distance < min) then
19:                 min = distance
20:                 AddressOfClosest = j + MaxNumInputs
21:             end if
22:         end if
23:     end for
24:     return AddressOfClosest
25: end function
```