Evolving Modular Genetic Regulatory Networks

Josh Bongard (bongard@ifi.unizh.ch)
Artificial Intelligence Laboratory
University of Zürich, Switzerland

Abstract—In this paper we introduce a system that combines ontogenetic development and artificial evolution to automatically design robots in a physics-based, virtual environment. Through lesion experiments on the evolved agents, we demonstrate that the evolved genetic regulatory networks from successful evolutionary runs are more modular than those obtained from unsuccessful runs.

I. Literature Review

The recent renaissance of 'evo-devo'[24]—evolutionary developmental biology—is causing a radical change in our understanding of how selection pressure shapes the organism's underlying genetic regulatory network (GRN).

Several startling discoveries have been made regarding the so-called Hox genes (master control genes that specify and order the body segments in most metazoan species[10]), including mounting evidence that these genes are highly conserved over many species[6][18], diversification of Hox gene clusters has led to a diversification in animal body plans[9][16], and that these genes are arranged along the chromosome in the same order that they are expressed along the anterior-posterior axis of the embryo[13]. However, insights into how selection pressure has shaped the evolution and diversification of such genes is only now beginning to appear in the literature[5][15].

In parallel to this, both neuroscience researchers and evolutionary biologists have postulated that modularity (integration of functionally related structures, and dissociation of unrelated structures) is necessary at both phenotypic[20][21] and genotypic[22][23] levels in order to evolve complex structures.

In the field of evolutionary robotics, evolutionary computation is now being used to evolve both the brains and bodies of virtual[19][1] and real-world robots[14], and focus is increasingly coming to bear on making the genetic encoding of these systems as modular and compact as possible in order to increase evolvability[11][4]. Eggenberger[8] first incorporated GRNs into an evolutionary simulation to evolve three-dimensional shapes. In this paper I report new results obtained from the Artificial Ontogeny system (AO), which grows virtual agents from GRNs and evaluates them in a physically-realistic, three-dimensional virtual environment[3].

II. Methods

Artificial Ontogeny extends the genetic algorithm to include ontogenetic development. In the results presented below, agents are tested for how fast they can travel over an infinite horizontal plane during a pre-specified time interval. The fitness determination is a two-stage process: the agent is first grown from a GRN (the growth phase), and then evaluated in its virtual environment (the evaluation phase) (Fig. 1).

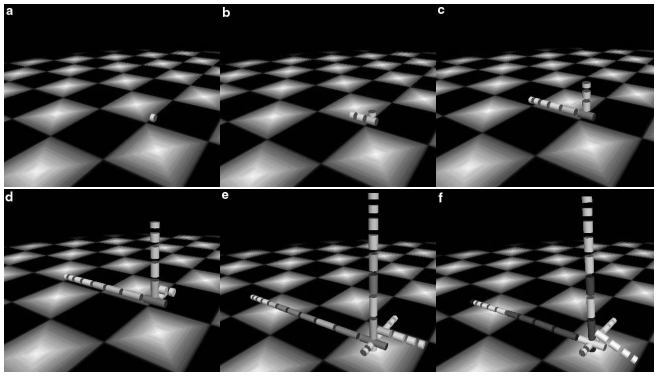

Figure 1: **a-e**: Images taken from t_0, t_{75}, t_{150}, t_{225} and t_{300} during the growth phase of an evolved agent. The units are darkened in proportion to how many neurons and synapses they contain. **f**: t_0 of the evaluation phase. The grey units contain motorized joints; the dark grey units are in contact with the ground plane.

Agents are composed of one or more cylindrical morphological units and zero or more sensors, motors, neurons and synapses. At the beginning of the growth phase, the genome to be tested and a motor neuron are inserted into a single unit. Two different transcription factors (TFs) are injected into the anterior and posterior poles of the unit, in order to allow the GRN to establish major body axes in the developing agent, if required (this has been shown to be one of the primary roles of maternal TF diffusion during early development[2]). The maternal TFs affect the expression of the zero or more genes lying along the genome embedded in the starting unit, which in turn may begin to emit TFs throughout the unit. The TFs may directly affect the phenotype of the developing agent: there are 23 pre-defined phenotypic transformations that TFs can initiate, such as increasing the length of a unit, causing a unit to split into two units, or adding, deleting or modifying the properties of the agent's neurons or synapses (Fig. 2).

Unlike the recursive parametric encoding schemes mentioned above, each genome in the AO system is treated as a genetic regulatory network[12][8][17], in which genes produce transcription factors that either have a direct phenotypic effect or regulate the expression of other genes.

Each genome to be evaluated is scanned by a parser, which marks the site of promotor sites. Promotor sites indicate the starting position of a gene along the genome, and are not hand-coded, but rather the number and position of them is under evolutionary control, similar to the

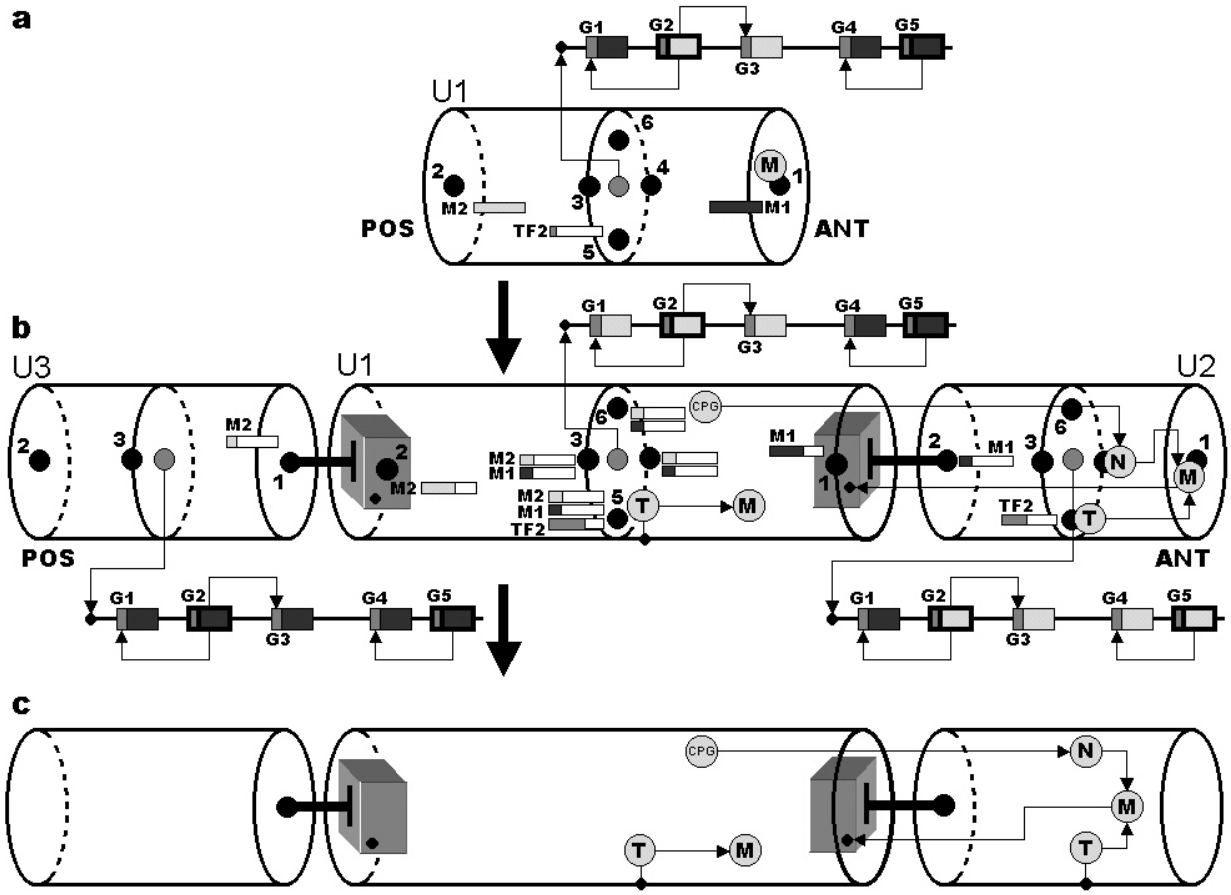

Figure 2: **a**: A hypothetical agent at the beginning of growth. The anterior direction (the direction the agent must move in order to gain fitness) is indicated (*ANT*), as is the posterior direction (*POS*). A genome, a motor neuron (*M*) and two maternal TFs (*M1*, *M2*) are injected into the single, beginning morphological unit (*U1*). The unit contains six TF diffusion sites (*1-6*). The genome contains five genes: *G1*, *G3*, *G4* are structural genes; *G2* and *G5* (outlined in bold) are regulatory genes. *G3* and *G5* are initially switched on, and begin to diffuse TFs into the unit; the other genes are initially switched off (light grey indicates expression; dark grey indicates repression). **b**: After several time steps, *U1* has split twice, producing neighbouring daughter units *U2* and *U3*, which are attached to it by one degree-of-freedom damped, torsional joints. The genome has been copied into *U2* and *U3*, where different combinations of TF concentrations have changed the states of some of the genes. *U1* has been lengthened by TF2, which increases unit length, released by *G3* at diffusion site 5. *M1* and *M2* have diffused throughout the unit. The motor neuron in *U1* has differentiated into a local neural circuit through combined gene action (*T*=touch sensor, *CPG*=central pattern generator, *N*=neuron). **c:** The fully grown agent from which all genetic material has been removed, in preparation for agent evaluation. The joint near *U2* is active, because it receives motor commands from the neural circuit in *U2*. The joint near *U3* is passive, and will swing freely during the evaluation phase because the motor neuron in *U3* has been deleted.

method employed in [17]. On average, there are 10 promotor sites, and thus 10 genes, found in any randomly generated genome.

Fig. 3 shows a magnification of gene *G3* from Fig. 2. The six floating-point values following a gene's promotor site supply the parameter values for the gene. The first value ($P1$) indicates which of the 20 possible TFs regulates the gene's expression. The second value ($P2$) indicates which of the 23 possible TFs is produced if this gene is expressed. The third value ($P3$) indicates which of the 6 TF diffusion sites the TF is diffused from if this gene is expressed. The fourth value ($P4$) indicates the concentration of the TF that should be injected into the diffusion site if the gene is expressed. The fifth and sixth values ($P5$ and $P6$) denote the concentration range of the regulating TF to which the gene responds.

All 43 TFs (23 TFs that directly affect the phenotype, and 20 regulatory TFs) share the same fixed, constant diffusion coefficients. For each time step that a gene emits its TF, the concentration of that TF, at the diffusion site encoded in the gene, is increased by the amount encoded in the gene (which ranges between 0.0 and 1.0), divided by 100. All TF concentrations, at all diffusion sites, decay by 0.005 at each time step. TFs diffuse between neighbouring diffusion sites within a unit at one-half this rate. TFs diffuse between neighbouring units at one-eighth the rate of intra-unit diffusion.

The agent's behaviour is dependent on the real-time propagation of sensory information through its neural network to motor neurons, which actuate the agent's joints.

There are two types of sensors that artificial evolution may embed within the units of the agent: touch sensors and

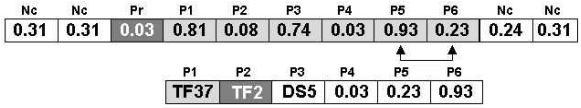

Figure 3: **A sample gene.** This gene (*G3* in Fig. 2) emits TF 2 from diffusion site 5 (*DS5*) if it is expressed (the concentration of TF 2 is increased by 0.03 at *DS5* during each time step of the growth phase that *G3* is expressed). If the average concentration of TF 37 in the current unit is between 0.23 and 0.93 the gene is expressed; otherwise, it is repressed. The gene is flanked by non-coding values (*Nc*).

proprioceptive sensors. Touch sensor neurons return a maximal positive signal if the unit in which they are embedded is in contact with either the target object or the ground, or a maximal negative signal otherwise. Proprioceptive sensors return a signal commensurate with the angle described by the two rigid connectors forming the rotational joint within that unit. The agent can also contain central pattern generator (CPG) neurons. These neurons emit a sinusoidal output signal: their frequency is modulated by the strength of the incoming signal (large positive input produces a high frequency, and large negative input produces a low frequency), and their phase is set relative to the time step (during the growth phase) when they are formed. Internal neurons can also be incorporated by evolution into an agent's neural network, in order to propagate signals from sensor to motor neurons. Finally, bias neurons emit a constant, maximum positive value.

The agent achieves motion by actuating its joints. This is accomplished by averaging the activations of all the motor neurons within each unit, and scaling the value between $-\frac{\pi}{2}$ and $\frac{\pi}{2}$ (these minimum and maximum joint angles may be reduced by the presence of one of the TFs that affects morphogenesis). Torque is then applied to the rotational joints such that the angle between the two rigid connectors forming the joint matches this value. The desired angle may not be achieved if: there is an external obstruction; the units attached to the rigid connectors experience opposing internal or external forces; or the values emitted by the motor neurons change over time. Note that failure to achieve the desired angle may be exploited by evolution, and may be a necessary dynamic of the agent's actions. If a unit contains no motor neurons, the rotational joint in that unit is passive.

II. Results

The agents reported in this section were evaluated in a three-dimensional, physically-realistic simulation package [1]. During each time step of the evaluation, sensor readings are taken, the neural network is updated, and the motor commands are translated into torques. The torques are passed to the simulator, which updates the positions, velocities and orientations of each of the agent's units. The updates are also affected by simulated external forces such as gravity, inertia, friction and collision or contact with the ground plane[2].

[1] Critical Mass Labs, `www.cm-labs.com`.

[2] By evaluating the agent in a physically realistic simulation, agents can evolve to take advantage of their environment, such as using gravity and momentum to move non-actuated joints in a useful manner. Also, it may be easier to translate evolved solutions into real-world robots.

Sixty independent evolutionary runs of 300 generations each were conducted, using a population size of 300. The initial population was composed of 300 strings of 200 floating-point values, rounded to two decimal places and ranging between 0.00 and 1.00. Genomes were evolved to maximize the fitness function

$$f = s + (p_z(t_{500}) - p_z(t_{250})) \sum_{t=1}^{500} \sum_{i=1}^{u_{tot}} |j_i(t)|, \quad (1)$$

$$s = n + m + sy + sy_{nz} + o_{nz}, \quad (2)$$

$$n = \begin{cases} u_{tot} & : \quad u_{tot} \leq 3 \\ 3 & : \quad u_{tot} > 3 \end{cases} \quad (3)$$

$$m = \begin{cases} 1 & : \quad s_{tot} > 0 \quad \text{and} \quad m_{tot} > 0 \\ 0 & : \quad \text{otherwise} \end{cases} \quad (4)$$

where u_{tot} is the number of units comprising the agent; $j_i(t)$ is the desired angle command sent to joint j in unit i at time step t; and $p_z(t_{500})$ and $p_z(t_{250})$ are the z-components of the anterior-most unit's position at the end of, and halfway through the evaluation period, respectively[3]. s is a shaping function: it awards agents that have not yet achieved any locomotion for particular phenotypes that favour the discovery of locomotion. n awards for agents that are composed of at least three units, and m awards for creatures that contain at least one sensor and one motor (s_{tot} and m_{tot} denote the total number of sensors and motor neurons in the agent, respectively). $sy = 1$, $sy_{nz} = 1$ and $o_{nz} = 1$ if the agent contains at least one synapse, one synapse with non-zero weight, or one non-zero motor neuron output, respectively, and are set to zero otherwise. The shaping function allows evolution to rapidly produce an agent that exhibits some active behaviour. An alternative approach would have been to seed evolution with minimally behaving agents, and omit the shaping function.

Strong elitism was employed; the best 150 genomes at each generation were retained. The mutation rate was set to produce, on average, random replacement of a single value for each new genome. Also, new genomes had a 10% chance of having a substring of their values excised (the length of the excised substring was chosen between 1 and $l - 1$ with a uniform distribution, where l is the length of the genome), and a 10% chance of two non-overlapping substrings (chosen between 1 and $\frac{l}{2} - 1$ with uniform distribution) from being swapped within the genome. Unequal crossover was employed, which allowed for gene duplication and deletion. Tournament selection, with a tournament size of 3, was used to select genomes to participate in crossover.

Fig. 1**f** shows the morphology of the most fit agent taken from one of the evolutionary runs; Figs. 4**a** and 5**b** show the morphologies of the most fit agents from two other runs.

In several of the runs, forward locomotion did not evolve; agents either exhibited random actuation, or discovered a way to fall over just after t_{250}. In other runs, small agents composed of no more than 6 units, and only 1 or 2 active joints, discovered forward locomotion. However in two runs, large agents (Figs. 1**f** and 4**a**) with several actuated joints achieved forward locomotion. The GRN of one of

[3] By ignoring any locomotion before t_{250}, agents that passively fall over receive low fitness values.

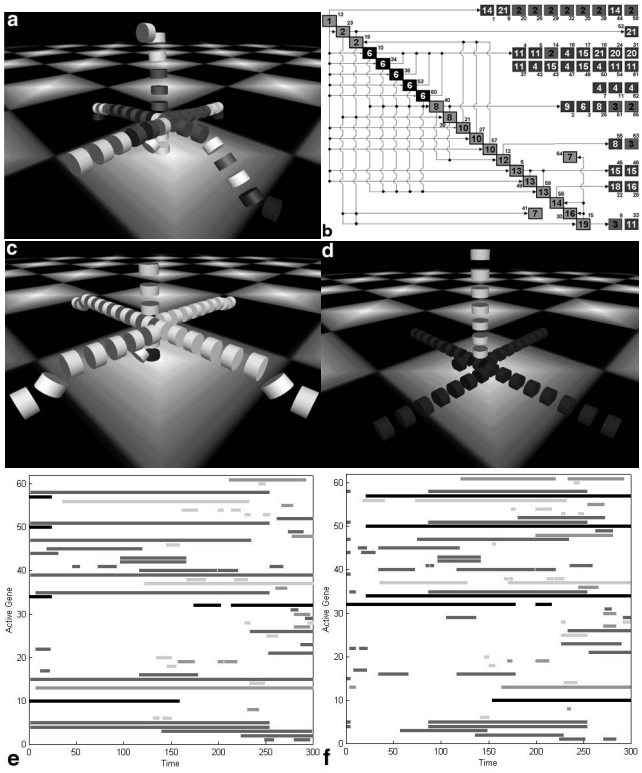

Figure 4: **Results from a lesion experiment. a**, The morphology of the most successful agent from one evolutionary run (wild-type). **b**, The underlying GRN specifying the agent's growth. **c**, The agent regrown with regulatory genes 10, 34, 35, 60 and 63 repressed in all units (loss-of-function). **d**, The agent regrown with the targetted genes expressed in all units (gain-of-function). **e**, Differences in gene expression between the first units of the wild-type and loss-of-function agents. Black bars indicate the targetted regulatory genes; dark grey bars indicate structural genes that influence neural growth; grey bars indicate structural genes that influence morphological growth; light grey bars indicate other regulatory genes. **f**, Differences in gene expression between the first units of the wild-type and gain-of-function agents.

these agents is shown in Fig. 4**b**. This genome contained 66 active genes—genes that were expressed for at least one time step, in at least one of the agent's units. Regulatory genes are depicted as boxes with bold edges; the other genes are structural genes. The black genes indicate the regulatory genes targetted for the lesion experiments shown in Fig. 4. The dark grey structural genes denote those genes that participate in neurogenesis: they guide the growth of the agent's neural structure. The grey structural genes participate in morphogenesis: they direct the growth of the agent's body. The numbers inside the gene indicate which TF is emitted by that gene. The numbers outside the genes indicate their relative position along the gene: gene 1 is the first active gene in the genome; gene 2 is the second, and so on. Arrows indicate gene regulation: for example, genes 45 and 46 are regulated by regulatory TF 13, which is emitted by genes 49 and 59. Genes 7, 11 and 62 are regulated directly by the anterior maternal TF. Genes 13 and 23 have evolved to emit the posterior maternal TF.

The set of genes that directly regulate the most neurogenesis genes (for this agent, genes 10, 34, 35, 60 and 63) were selected for mutation. The agent was regrown with these genes suppressed in all units: Fig. 4**c** shows the morphology of this loss-of-function mutant. The agent was then regrown again, with these genes expressed in all units: Fig. 4**d** shows the morphology of this gain-of-function mutant. The expression pattern differences of the 66 active genes between the first unit of the original (wild-type) agent[4] and the loss-of-function mutant are shown in Fig. 4**e**. The expression pattern differences of the 66 active genes between the first unit of the original agent and the gain-of-function mutant are shown in Fig. 4**f**.

III. Analysis

As can be seen from Figs. 4**e** and 4**f**, the supression or enhancement of the five targetted regulatory genes has a larger effect on the structural neurogenesis genes than that on the morphogenesis genes. Similarly, the morphologies of the loss-of-function and gain-of-function mutants (Figs. 4**c** and 4**d**) are quite similar to that of the wild-type agent. However, in both agents, the neural disruption was severe enough such that none of the joints were actuated. This indicates that there is high pleitropy (co-regulation) between the neurogenesis genes, and lower pleitropy between neurogenesis and morphogenesis genes. In other words, a dissociation between regulation of neurogenesis and morphogenesis has occurred: that is, evolution can experiment with different body plans and not disrupt neurogenesis, and can experiment with different neural components on the same body plan.

A measure has been formulated to quantify this genetic modularity, using the weighted sums

$$N_W^L = \frac{\sum_{u=1}^{u_{tot}} t^{(u)} \sum_{t=1}^{300} \sum_{i=1}^{g_n} |g_i^W(t) - g_i^L(t)|}{\sum_{u=1}^{u_{tot}} t^{(u)}} \quad (5)$$

$$M_W^L = \frac{\sum_{u=1}^{u_{tot}} t^{(u)} \sum_{t=1}^{300} \sum_{i=1}^{g_m} |g_i^W(t) - g_i^L(t)|}{\sum_{u=1}^{u_{tot}} t^{(u)}} \quad (6)$$

$$N_W^G = \frac{\sum_{u=1}^{u_{tot}} t^{(u)} \sum_{t=1}^{300} \sum_{i=1}^{g_n} |g_i^W(t) - g_i^G(t)|}{\sum_{u=1}^{u_{tot}} t^{(u)}} \quad (7)$$

$$M_W^G = \frac{\sum_{u=1}^{u_{tot}} t^{(u)} \sum_{t=1}^{300} \sum_{i=1}^{g_m} |g_i^W(t) - g_i^G(t)|}{\sum_{u=1}^{u_{tot}} t^{(u)}} \quad (8)$$

where N_W^L and N_W^G indicate the expression differences between the neurogenesis genes in the wild-type agent and loss-of-function mutant, and the wild-type agent and gain-of-function mutant, respectively. A value of zero indicates there were no expression differences between any of the neurogenesis genes; a value of one indicates that whenever a neurogenesis gene—at any time step in any unit of the wild-type agent—is expressed (or suppressed), it is suppressed (or expressed) during that time step, in that unit, of the mutant. Similarly, M_W^L and M_W^G indicate the expression differences between the morphogenesis genes in the wild-type agent and loss-of-function mutant, and the wild-type agent and gain-of-function mutant, respectively.

[4]The term 'wild-type' refers to the fact that the agent was grown from a genome taken directly from a completed evolutionary run, and no additional modifications have yet been made to it.

u_{tot} here indicates the total number of units comprising the wild-type, loss-of-function or gain-of-function agent with the minimum number of units. g_n and g_m indicate the number of active neurogenesis and morphogenesis genes in the wild-type agent, respectively. $g_i^W(t) > 0$, $g_i^L(t) > 0$ and $g_i^G(t) > 0$ if gene i in the wild-type, loss-of-function or gain-of-function agent is expressed at time step t; and are set to zero otherwise. $t^{(u)}$ indicates the number of time steps for which unit u is present during the growth phase: $t^{(1)} = 300$, and units appearing later during the growth phase have lower values.

In some agents, suppressing or enhancing the targetted regulatory genes disrupts the morphology such that the second and subsequent units in the loss-of-function or gain-of-function mutants appear earlier or later than they do in the wild-type agent. Thus, in order to compare the expression patterns of genes between these units, the expression patterns are expanded from binary strings with lengths less than 300 to floating-point strings of length 300 with values in $[0, 1]$ using bilinear scaling[7].

Now, the pairs $[N_W^L, M_W^L]$ and $[N_W^G, M_W^G]$ indicate the relative neurological and morphological effects caused by artificially suppressing or enhancing the expression of targetted regulatory genes. This measure was applied to the most fit agent from each of the 60 runs; the targetted gene set was chosen by selecting those regulatory genes that directly co-regulated the maximum number of neurogenesis genes. In some agents, the loss-of-function mutation had a greater effect than the gain-of-function mutation, and in other agents, the reverse case was true, depending on how the targetted genes are expressed in the wild-type agent. In order to compare mutational effect, if $N_W^L > N_W^G$ in an agent, then $[N_W^L, M_W^L]$ was retained and $[N_W^G, M_W^G]$ was discarded; otherwise, $[N_W^G, M_W^G]$ was retained and $[N_W^L, M_W^L]$ was discarded. Fig. 5**a** plots these 60 remaining value pairs: it shows the relative neurological versus morphological effects of the lesion experiment on each agent.

As can be seen, the two runs that produced the large, locomoting agents produced more highly modular GRNs that the GRNs evolved in the other evolutionary runs: lesioning of the targetted genes in these agents had quite a drastic neurological effect, but a relatively mild morphological effect. Moreover, the agent with the most modular GRN had the maximum number of actuated joints, indicating a relatively sophisticated neural architecture (see Fig. 5, inset), even though it did not exhibit much forward locomotion. In addition, the evolutionary history of these three runs was searched, and the agent in which the targetted genes appeared were located. These three agents were then lesioned as well, and it was found that in all three runs, the targetted gene had no morphological effect at all. This suggests that part of the reason for the evolutionary success of these populations is due to the early appearance of highly modular GRNs.

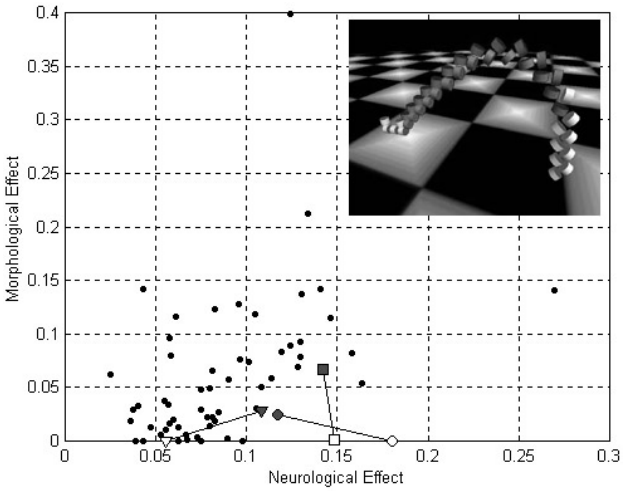

Figure 5: **Plot of neurological versus morphological effect from 60 lesion experiments.** The filled triangle, square and circle correspond to the agents shown in Figs. 1**f**, 4**a** and (inset). The open triangle, square and circle correspond to the first agents appearing in these three evolutionary runs that contained the targetted regulatory gene. **(inset)**: The evolved agent with the most actuated joints.

IV. Conclusions

In this paper we have outlined the workings of the Artificial Ontogeny system (AO), which incorporated ontogenetic development into the artificial evolution of behaving agents. It has been demonstrated that this system can be used to evolve locomoting agents with a high part count. Finally, it was shown that part of the reason for the evolutionary success of these populations was due to the early evolution of modular genetic regulatory networks: the genomes exhibited high pleitropy between the genes responsible for neural growth, and low pleitropy between the genes responsible for neural and morphological growth.

Because this system acts as an abstract model of both evolution and development, it is extremely general. It can be used to test several hypothesis about how adaptive changes to the developmental programme of an evolving population is affected by behavioural selection pressure. To the best of our knowledge, this paper has provided for the first time quantitative data on how behavioural selection pressure shapes genetic regulatory networks. Moreover, the large neurological effects exhibited by the regulatory genes in the successful evolutionary runs indicates that these genes are acting like master control genes. This indicates that the AO system may be very useful for testing hypotheses about how *Hox* genes have evolved in nature.

Future studies are planned for directly comparing the phenotypes of wild-type agents and lesioned mutants, in order to clarify how phenotypic and genotype modularity are related. Also, experiments are planned with the AO system for investigating how and why some regulatory genes come to adopt a master control role during development.

References

[1] A. Adamatzky, M. Komosinski & S. Ulatowski, "Software review: Framsticks", Kybernetes: The International Journal of Systems & Cybernetics, Vol. 29, pp. 1344–1351, 2000.

[2] K. V. Anderson, & C. Nüsslein-Vollhard, "Information for the dorso-ventral pattern of the Drosophila embryo

is stored in maternal mRNA", Nature, Vol. 311, pp. 223–227, 1984.

[3] J. Bongard & R. Pfeifer, "Repeated Structure and Dissociation of Genotypic and Phenotypic Complexity in Artificial Ontogeny", Proceedings of The Genetic and Evolutionary Computation Conference, GECCO-2001, pp. 829–836, 2001.

[4] R. Calabretta, S. Nolfi, D. Parisi & G. P. Wagner, "Duplication of modules facilitates the evolution of functional specialization", Artificial Life, Vol. 6, pp. 69–84, 2000.

[5] S. B. Carroll, "Endless forms: the evolution of gene regulation and morphological diversity", Cell, Vol. 101, pp. 577–580, 2000.

[6] J. Cohen, "Development of the zootype", Nature, Vol. 363, p. 307, 1993.

[7] T. F. Cox, M. A. A. Cox, Multidimensional Scaling, CRC Press, 2000.

[8] P. Eggenberger, "Evolving morphologies of simulated 3D organisms based on differential gene expression", Procs. of the Fourth European Conference on Artificial Life, pp. 205–213, 1997.

[9] J. R. Finnerty, "Head start", Nature, Vol. 408, p. 778, 2000.

[10] W. Gehring, Master Control Genes in Development and Evolution, Yale University Press, 1997.

[11] F. Gruau, "Automatic definition of modular neural networks", Adaptive Behaviour, Vol. 3, pp. 151–183, 1994.

[12] S. A. Kauffman, The Origins of Order, Oxford University Press, 1993.

[13] E. B. Lewis, "Clusters of master control genes regulate the development of higher organisms", Journal of the American Medical Association, Vol. 267, pp. 1524–1531, 1992.

[14] H. Lipson, & J. B. Pollack, "Automatic design and manufacture of robotic lifeforms", Nature, Vol. 406, pp. 974–978, 2000.

[15] R. S. Mann, "Why are Hox genes clustered?", Bioessays, Vol. 19, pp. 661–664, 1997.

[16] A. Meyer, "Hox gene variation and evolution", Nature, Vol. 391, p. 225, 1998.

[17] T. Reil, "Dynamics of gene expression in an artificial genome—implications for biological and artificial ontogeny", Proceedings of the Fifth European Conference on Artificial Life, pp. 457–466, 1999.

[18] B. Schierwater & R. Desalle, "Current problems with the zootype and the early evolution of Hox genes", Mol. Dev. Evol., Vol. 291, pp. 169–174, 2001.

[19] K. Sims, "Evolving 3D morphology and behaviour by competition", Procs. of Artificial Life IV, pp. 28–39, 1994.

[20] G. Tononi, O. Sporns, & G. M. Edelman, "A measure for brain complexity: Relating functional segregation and integration in the nervous system", Proc. Natl. Acad. Sci. USA, Vol. 91, pp. 5033–5037, 1994.

[21] G. Tononi, O. Sporns, & G. M. Edelman, "Measures of degeneracy and redundancy in biological networks", Proc. Natl. Acad. Sci. USA, Vol. 96, pp. 3257–3262, 1999.

[22] G. Wagner, "Homologues, natural kinds and the evolution of modularity", Amer. Zool., Vol. 36, pp. 36–43, 1996.

[23] G. Wagner, & L. Altenberg, "Complex adaptations and the evolution of evolvability", Evolution, Vol. 50, pp. 967–976, 1996.

[24] K. P. White, "Functional genomics and the study of development, variation and evolution", *Nature Reviews Genetics* Vol. 2, pp. 528–537, 2001.