

The Triangle of Life: Evolving Robots in Real-time and Real-space

A.E. Eiben¹, N. Bredeche^{2,3}, M. Hoogendoorn¹, J. Stradner⁴, J. Timmis⁵, A.M. Tyrrell⁶, A. Winfield⁷

¹VU University Amsterdam (a.e.eiben@vu.nl, m.hoogendoorn@vu.nl)

²UPMC Univ Paris 06, UMR 7222, ISIR, F-75005, Paris, France (nicolas.bredeche@upmc.fr)

³CNRS, UMR 7222, ISIR, F-75005, Paris, France

⁴Karl-Franzens University Graz (juergen.stradner@uni-graz.at)

⁵University of York (jon.timmis, andy.tyrrell@york.ac.uk)

⁶Bristol Robotics Lab, (alan.winfield@uwe.ac.uk)

Abstract

Evolutionary robotics is heading towards fully embodied evolution in real-time and real-space. In this paper we introduce the Triangle of Life, a generic conceptual framework for such systems in which robots can actually reproduce. This framework can be instantiated with different hardware approaches and different reproduction mechanisms, but in all cases the system revolves around the conception of a new robot organism. The other components of the Triangle capture the principal stages of such a system; the Triangle as a whole serves as a guide for realizing this anticipated breakthrough and building systems where robot morphologies and controllers can evolve in real-time and real-space. After discussing this framework and the corresponding vision, we present a case study using the SYMBRION research project that realized some fragments of such a system in modular robot hardware.

Introduction

Evolutionary robotics is heading towards fully embodied evolution in real-time and real-space. In this paper we introduce the Triangle of Life, a general conceptual framework that can help build systems where robots can actually reproduce. The framework can be instantiated with different hardware approaches and different reproduction mechanisms. For example, one could use classic mechatronic components and 3D-printing to produce new robots, or a stock of autonomous actuated robot modules as raw material and self-driven aggregation to implement ‘birth’.

The novelty of this framework lies in the pivotal role of reproduction and conception. The life cycle it captures does not run from birth to death, but from conception to conception and it is repeated *in real hardware* thus creating ‘robot children’ over and over again. This is new in evolved 3D printed robots, where the body structure is printed off-line. Even if the design is evolved, the printer only produces the end result after evolution is halted (in simulation), whereas in our framework printing=birth, thus being part of the evolutionary process, rather than following it.

Our approach is also new in self-assembling robot swarms, because existing work traditionally focusses on the transition of a swarm into an aggregated structure (a robot

organism) and vice versa. In the traditional setting, being aggregated is a transient state that enables the robots to meet a certain challenge after which they can disassemble and return to normal. In contrast, we perceive being aggregated as a permanent state and consider aggregated structures as viable robotic organisms with the *ability to reproduce*. That is, two or more organisms can recombine the (genetic) code that specifies their makeup and initiate the creation of a new robotic organism. This differs from earlier work aiming at *self-replication* and *self-reconfiguration* in that a ‘child organism’ is neither a replica of its parents, nor is it a reconfigured version of one of them.

This paper has a twofold objective, 1) to present the Triangle of Life as a conceptual framework for creating ALife of this type and 2) to illustrate how the components of this framework can be implemented in practice. To this end, we will use the SYMBRION research project¹ as a case study, even though originally the project only targeted traditional swarm-to-organism-to-swarm systems, cf. Levi and Kernbach (2010).

Background and related work

The ideas in this paper can be considered from three perspectives, that of artificial life, evolutionary computing, and (evolutionary) robotics. The modern scientific vision of creating artificial life has a long history dating back to the 1987 Santa Fe workshop, cf. Langdon (1989); Levy (1992); Langton (1995). The most prominent streams in the development of the field are traditionally based on wetware (biology and/or chemistry), software (i.e., computer simulations), and hardware (that is, robots). In this paper we focus on the third option. The main contribution of the paper from this perspective is the introduction of a new, integrative framework, the Triangle of Life, that helps develop and study hardware-based ALife systems. In fact, the Triangle of Life defines a new category of ALife systems and outlines an interesting avenue for future research.

¹EU Grant number FP7-ICT-2007.8.2, running between 2008-2013.

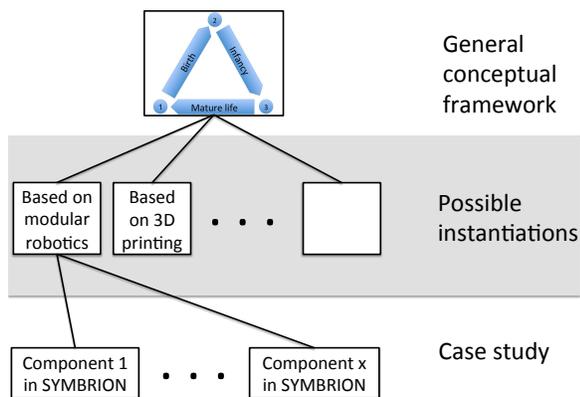


Figure 1: Positioning the Triangle of Life, its possible instantiations in general, and the specific examples used in this paper.

From an evolutionary perspective the framework we advocate here corresponds to a major transition from evolutionary computing (i.e., artificial evolution in software) to Embodied Artificial Evolution (i.e., artificial evolution in hardware) as introduced in Eiben et al. (2012). The roadmap outlined there considers embodiment in the broad sense, including biochemical approaches and treats mechatronics based embodied evolution as one of the possible incarnations. The work presented here represents the first detailed elaboration entirely devoted to that kind of systems.

Finally, the vision behind this paper can also be considered from the perspective of robotics. The relevant subarea here is evolutionary robotics that has a large body of related work, e.g., Nolfi and Floreano (2000); Wang et al. (2006); Trianni (2008). However, most existing systems in this field are based on simulations and use evolutionary algorithms as optimizers in an off-line fashion, during design time. Furthermore, evolution is usually applied to optimize/design some parts of the robot morphology or the controller, but rarely both of them. In contrast, our vision concerns real hardware, on-line evolution during run time, and it includes the evolution of both the morphologies and the controllers. In the system we envision, new robots are produced continuously only limited by the availability of the raw materials and the capacity of the ‘birth’ mechanism. In the resulting system evolution is not a simple optimizer of some robot features, but a force of continuous and pervasive adaptation.

In the landmark Golem project Lipson and Pollack (2000) evolved robots capable of moving themselves across a flat surface; robots were evolved in simulation and the fittest individuals then fabricated by first 3D printing the structural components then adding motors to actuate the robot. Although a remarkable achievement, the artificial creatures evolved then physically realized contained neither sens-

ing nor controller, so were not self-contained autonomous robots. Only the robot’s physical morphology was evolved.

The use of Lego has featured in evolutionary robot hardware. Although not evolving complete Lego robots work has been described, and indeed attempted to formalise the use of Lego structures for evolution. For example Funes and Pollack (1997) describe the simulated evolution, then construction using Lego, of physical bridge-like structures. Peysakhov et al. (2000) present a graph grammar for representing and evolving Lego assemblies, and Devert et al. (2006) describe BlindBuilder, an encoding scheme for evolving Lego-like structures.

Notably Lund (2003) describes the “Building Brains and Bodies approach” and demonstrates the co-evolution of a Lego robot body and its controller in which the evolved robot is physically constructed and tested. Here simulated evolution explores a robot body space with 3 different wheel types, 25 possible wheel positions and 11 sensor positions. Lund observes that although the body search space is small, with 825 possible solutions, the search space is actually much larger when taking into account the co-evolved controller parameters. This work is significant because it is, to the best of our knowledge, the only example to date of the simulated co-evolution, then physical realisation, of body morphology and controller for a complete autonomous mobile robot.

Work by Zykov et al. (2007) describes an evolving modular robotic system on the Molecube platform. In this work, self-reproduction is not a necessary prerequisite of evolution, but rather its target. In particular, the authors evolve self-replicators by employing a genetic algorithm (in a 2D simulation) where the measured amount of self-replication is used as an explicit fitness criterion to evaluate morphologies. Then, in a second stage they evolve a command sequence, i.e., controller, that enables a given morphology to produce an identical copy of itself. However, as yet, there is still no work that has fully demonstrated the online evolution of both structure and function of a modular robotic system, that is fully embodied in the modules themselves.

A related area with practical relevance to our vision is that of self-organizing robotic systems, Murata and Kurokawa (2012). Modular self-reconfigurable robot systems, cf. Yim et al. (2007), are particularly interesting because they constitute one of the possible technologies for implementing the Triangle of Life as shown in Figure 1. However, conceptually such systems are quite different from ours, because the emphasis is on self-reconfiguring morphologies to adapt to dynamic environments, whereas in our evolutionary system, new morphologies appear through ‘birth’ and adaptation of morphologies takes place over generations.

The Triangle of Life

Throughout this paper we will not attempt to (re)define what life is. Instead, we take a pragmatic approach and con-

sider three features that are typically attributed to life or life-like systems: self-reproduction that relies on heredity, self-repair, and learning.

The proverbial Cycle of Life revolves around birth. We adopt this stance and define the Triangle of Life as shown in Figure 2.

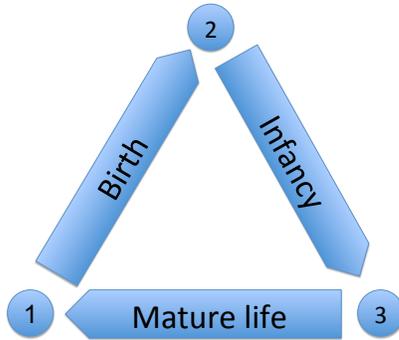


Figure 2: The Triangle of Life. The pivotal moments that span the triangle are: 1) Conception: A new genome is activated, construction of a new organism starts. 2) Delivery: Construction of the new organism is completed. 3) Fertility: The organism becomes ready to conceive offspring.

This concept of the Triangle is generic, the only significant assumption we maintain is the genotype-phenotype dichotomy. That is, we presume that the robotic organisms as observed ‘in the wild’ are the phenotypes encoded by their genotypes. In other words, any robotic organism can be seen as the expression of a piece of code that we call the genome. As part of this assumption we postulate that reproduction takes place at the genotypic level. This means that the evolutionary operators mutation and crossover are applied to the genotypes (to the code) and not to the phenotypes (to the robotic organisms). This fundamental assumption not only makes our envisioned systems more life-like, but –perhaps even more importantly– keeps the door open to enhancing the system with developmental abilities.

In the forthcoming subsections we will elaborate on each stage of the Triangle. For the sake of clarity we appeal to the modular robotic approach and explain some details in that setting. However, we emphasize that the Triangle is a generic framework equally applicable to modular and non-modular approaches.

Birth

A new robotic organism is created first at genotype level and is thus seeded by a new piece of genetic code that is created by mutating or recombining existing pieces of code. Birth is therefore the first stage of life, specified as the interval between the moment of activating a newly created genome (circle 1 in Figure 2) and the moment when the robot organism encoded by this genome is completed (circle 2 in

Figure 2). In technical terms, this is the period when morphogenesis takes place. In principle, it can be implemented in various ways and later on we will illustrate some in detail. Here we suffice to distinguish two main categories, based on explicit vs. implicit representations of the shape of the newborn robot organism. Using an explicit representation, the genome explicitly specifies the shape of the organism and the process of morphogenesis is executed with this shape as target. Morphogenesis has therefore a clear stopping criterion; it is successfully completed when the target shape has been constructed. Using implicit representation the genome does not contain an exact description of the new shape. Rather, the genome can be seen as a set of rules governing the morphogenesis process that could follow different tracks and thus deliver different end shapes depending on the given circumstances and random effects. Note that this notion of implicit representation includes indirect, developmental representations, EvoDevo, ect. and connects our vision with the nascent area of morphogenetic engineering, cf. Doursat et al. (2012).

Infancy

The second stage in the Triangle of Life starts when the morphogenesis of a new robot organism is completed (circle 2 in Figure 2) and ends when this organism acquires the skills necessary for living in the given world and becomes capable of conceiving offspring (circle 3 in Figure 2). This moment of becoming fertile is less easy to define in general than the other two nodes of the triangle. However, we believe it is useful to distinguish an Infancy period for two reasons. Firstly, the new organism needs some fine tuning. Even though its parents had well matching bodies and minds (i.e., shapes and controllers), recombination and mutation can shuffle the parental genotypes such that the resulting body and mind will not fit well. Not unlike a newborn calf the new organism needs to learn to control its own body. Depending on the given system design this could take place under protected circumstances, under parental supervision or within an artificial ‘nursery’ with a food rich environment, etc. From this perspective, the Infancy interval serves as a grace period that allows the new organism to reach its full potential. Secondly, the new organism needs to prove its viability. System resources are expensive, thus should be allocated to the creation of offspring with an expectedly high quality. Introducing a moment of becoming fertile (after birth) implies that organisms must reach a certain age before they can reproduce. From this perspective, the Infancy period serves as an initial assessment of implicit fitness that helps filter out inferior organisms before they start wasting resources by producing offspring.

The moment of becoming fertile can be specified by any user-defined criterion. This could be as simple as time elapsed after birth, or some measurable performance, for instance, speed (high is good) or amount of energy collected

(large is good) or number of collisions with obstacles (low is good), etc.

Mature life

The third stage in the Triangle is the period of maturity. It starts when the organism in question becomes fertile (circle 3 in Figure 2) and leads to a new Triangle when this organism conceives a child, i.e., produces a new genome through recombination and/or mutation (circle 1).² It should be noted that at this point we switch perspectives: the beginning of a new life marks the beginning of another Triangle belonging to the new organism encoded by the new piece of genome. As for the ‘old’ organism nothing needs to end here. In other words, conceiving a child does not mean the end (death) of this organism, and it is certainly possible that an organism produces multiple offspring during its mature life. This view is motivated by the intuition behind the proverbial Cycle of Life that inspired our Triangle.

Robotic organisms can exhibit several behaviors during the mature period, depending on the given system and the interests of the experimenter. Here we will only consider two that we consider essential to any real world ALife system: reproduction and self-repair. Reproduction is an obvious requirement, but implementing it is challenging. For multi-cellular robotic organisms we see three feasible options:

1. Based on a ‘birth clinic’. After recombining the genomes of two parent organisms, the genome describing the new organism is beamed to a central facility where there are free robot modules. This is the place where the birth process is executed and a child robot is constructed.
2. Based on self-sacrifice. After recombining the genomes of two parent organisms, one of the parents disassembles and the child is built from its modules. Leftover modules become free riders and serve as available raw material. If the number of modules in the parent is not enough, others are recruited from such free riders.
3. A protocol based on seeds/eggs. This will be discussed later in detail as the one applied in SYMBRION.

Further to reproduction, we consider self-repair as an essential feature here. In simulation based ALife systems the world and its inhabitants can be stable and error-free, where randomness needs to be added deliberately. In the real-world systems we envision this is not the case, real hardware always breaks down. Thus, some form of self-repair is needed for continued operation after the inevitable breakdowns of the robot/organism. The ability to self-repair is linked to the

²Strictly speaking, the moment of producing a new genome need not be the same as activating this genome and starting the morphogenesis process, but this is just a formal detail with no real effect on the conceptual framework.

ability of the organism to perform morphogenesis, as it is very likely that some form of reconfiguration is needed in the event of failure.

Implementing the Triangle of Life

As mentioned in the Introduction, originally the SYMBRION project considered robotic organisms as transient states of the system. An aggregated organism could achieve goals a simple swarm could not (negotiating an obstacle or reaching a power point) and after completion it could dis-aggregate again. However after five years of research and development many of the components that make up the Triangle of Life have been implemented in hardware or are very close to being implemented in the short term. The purpose of this section is to illustrate these achievements together and to indicate the current state of the art towards an integrated ALife system based on the modular robotic organisms concept.

Birth: Explicit Encoding for Morphogenesis

Within the Symbion framework a heterogeneous group of mobile robots can operate in *swarm* mode to – for instance – autonomously explore a region, exploiting the spatial distribution of the swarm. However, when required, Symbion robots can self-assemble to form a 3D *organism*. The process of transition from swarm-mode to organism-mode, with an explicit pre-defined (or pre-evolved) body plan, is also self-organising and proceeds as follows. Any individual robot in swarm mode can act as a ‘seed’ robot, initiating morphogenesis. Typically this might be when that robot discovers an environmental feature or hazard that cannot be accessed or overcome by individual swarm-mode robots. Each robot is pre-loaded with a set of body-plans, and the seed robot will select the most appropriate body plan for the current situation. The position of the seed robot in the selected body plan then determines the next robot(s) that need to be recruited by the seed robot, and the face(s) that they will need to dock into. The seed robot then broadcasts message bearing recruitment signals from the selected face(s), using the IR signalling system built into each docking face. That message specifies which of the three Symbion robot types needs to be recruited.

The autonomous docking approach is illustrated in Figure 3. Initially, a seed robot initiates recruitment of other robots. The pre-evolved body plan is then transferred from the seed robot to them, so newly recruited robots then determine their own position in the growing organism. In discovering its position a robot also determines whether or not other robot(s) need to be recruited. In Figure 3 image 2 they do. Robots’ recruitment signals can be detected by other robots within range (150 cm) to provide rough directional information to any robots in range. IR beacon signals are used at short range (15 cm) to guide the approaching robots for precise alignment with the docking face. Upon com-

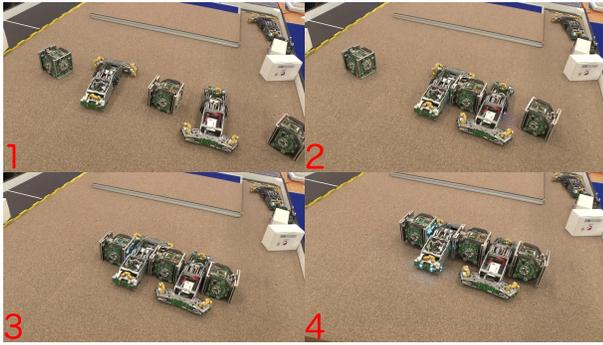


Figure 3: Morphogenesis in progress. Image 1: Five robots are in swarm mode. Image 2: Self-assembly is in progress. Image 3: The new organism is complete, but in 2D planar form. Image 4: The organism ‘stands up’ to transform to 3D.

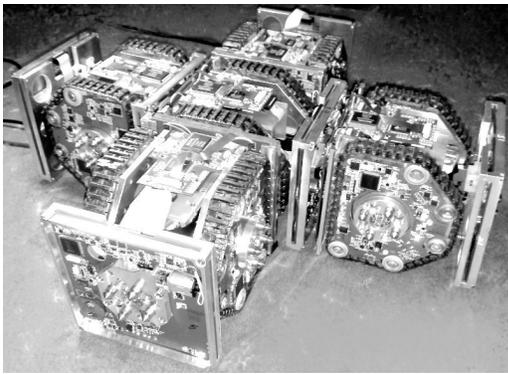


Figure 4: Example of a result from embodied morphogenesis using 5 Symbion robots obtained with the Virtual Embryogeny approach (credits: Markus Dauschan). See Dauschan et al. (2011) for details.

pletion of the docking process, robots stop emitting beacon signals. The same process is then repeated until the pre-evolved structure is formed. A behaviour-based approach is adopted for the design of the morphogenesis controller, together with a well-formatted tree structure which explicitly represents the organism body-plan, as described in Liu and Winfield (2012).

In this way robots initially form a 2D planar structure, see Figure 3 image 3. Once the robots in the 2D planar structure have assumed the correct functionality, according to their position in the body plan, the ‘organism’ will lift itself from 2D planar configuration to 3D configuration (as shown in Figure 3 image 4) and, with respect to locomotion, function as a macroscopic whole.

Birth: Implicit Encoding for Morphogenesis

An alternative to direct encoding is to consider developmental and generative systems (or *implicit* encodings). In

this setup, the information contained in the genome encodes the *process* of construction rather than an explicitly formulated *plan* of construction. While developmental and generative systems have been studied for some time (cf. the works of Bentley and Kumar (1999); Stanley and Miikkulainen (2003); Bongard and Pfeifer (2003)), the very process of morphogenesis starting from a swarm of autonomous units and going towards a full assembled organism raise additional issues, as the actual morphogenesis should be considered as an embodied process: online and decentralized.

In the last five years, several approaches have been investigated in the Symbion project, from theoretical ideas to practical robotic implementations, as shown in Figure 4. These approaches have been explored and tested, either with simulated or real robots, and have investigated the benefits of deterministic vs. stochastic morphogenesis from different perspectives (either bio-inspired or completely artificial). On one side, genetic regulatory networks (GRN) and artificial ontogenic process have been considered (Thenius et al. (2010)). On the other side, cellular automata (CA) have been used to model the developmental process by considering each robot of the organism as a cell with a von Neumann neighbourhood. In both cases, cells would be considered as homogeneous, that is sharing the same evolved update rules, whether this was explicit CA rules, a GRN update network and any other kind of developmental program. However, each cell would then trigger the recruitment of other cells depending on their current (possibly unique) situation, ultimately leading to a full-grown organism having reached a stable final configuration, as explored by Devert et al. (2011).

What makes these approaches particular with respect to the literature is that it is not only necessary to encode the morphogenesis process itself (i.e. the assembling sequence), but it is also mandatory to consider the actual execution of this process (the *embodied* morphogenesis): individual units are indeed facing a possibly challenging coordination tasks, with the possible constraints of satisfying temporal and spatial constraints (e.g. the assembly ordering can be important). Moreover, open-ended evolution of embodied morphogenesis can benefit from a creative process, that is to come up with original morphological solutions to address the challenges in the environment at hand. We devised a set of performance indicators to encompass these various desired properties and these are described below.

Evolvability is considered as the ability for the algorithm to produce *viable* shapes during the course of evolution. It is evaluated by counting the number of unique viable shapes out of a predefined number of tries.

Initial viability provides an indicator to estimate how difficult it is to bootstrap an evolutionary process. It is computed by considering only random generations of genotypic description for the encoding under scrutiny, and counting the number of shapes that can actually be build (i.e. viable)

out of the total number of shape descriptions generated.

Self-repair stands as one of the typical benchmark for morphogenesis and evaluates how a full organism can be successfully reconstructed from a starting condition that may not match the original initial condition (e.g. from the last recruited robot rather than from the original "egg" robot).

Lastly, *controllability* (unsurprisingly) evaluates the efficiency with respect to evolving the construction process to achieve a particular target shape: the faster the evolution, the better the controllability.

Infancy: Gait Learning

In our vision of Artificial Life based on hardware birth is followed by the stage of infancy. From an evolutionary point of view the proof of viability at the very beginning of this stage does not need any further consideration. If an organism, for example, consumes too much energy, its genome will not spread. Thus, in SYMBRION we concentrate on the objective of an organism learning to control its own body for locomotion. This is because movement increases the chances to spread the genome during the upcoming phase of mature life, independent of the chosen reproduction implementation during the mature phase. Thus, the objective of gait learning is an indirect one. The obvious easy solution of so-called free-riders, which are organisms staying in place and waiting for others to come by, can only exist in a low number in the population from an evolutionary perspective.

Here, gait learning comes with challenge of an unknown body shape. There may have been passed on a genome performing good locomotion from the ancestor but this good performance does not automatically hold for a different body shape. Thus, investigations on gait learning for a modular multi-robot organism –as it is the case in SYMBRION– always start from scratch.

As mentioned above on-line, on-board evolution was chosen in SYMBRION to be the optimization process. This leads to several important consideration and scientific questions. For example, the part of the genome which is responsible for locomotion could use Lamarckism. This means that at the beginning of mature life not the original genome but the genome altered by artificial evolution during gait learning is used for recombination.

The way of achieving shared control is another consideration. Should the controllers of the single modules in the organism be derived from an identical genome ("homogeneous")? Different genomes ("heterogeneous") would ease the creation of division of labor as some cells would be used to push, others to pull. In Waibel et al. (2009) it is stated that a homogeneous genome of team members is better suited when the task requires high level of cooperation.

Another important aspect is the type of controller being used. This strongly depends on the actuators used for locomotion. The three robot platforms which are the modules of



Figure 5: Multi-robot organism consisting of three modules during infancy. Screenshots show attempts of the organism to create locomotion during on-line, on-board evolutionary.

the multi-robot organism in SYMBRION come with several 2D actuators and one 3D actuator. The primary focus has been on the 3D drive. It is implemented as hinges which makes it possible to lift the other modules. Fig. 5 shows an example of the resulting 3D locomotion of an organism. This leads to a snake- or caterpillar-like motion. Three different controller types known for their evolvability were taken into consideration: CPG (central pattern generator), AHHS (artificial homeostatic hormone system, see Stradner et al. (2012)) and GRN (gene regulatory network). The idea is not to limit the population to one solution in the first place but to let evolution decide. The organism will only be controlled by one type during infancy phase, but the better it performs the greater the chance that this type will also be used by its offspring.

The ongoing work in SYMBRION is the implementation and testing (first results are shown in Fig. 5) for experiments to investigate the considerations raised above concerning gait learning for multi-robot organisms.

Mature Life: Self-Reproduction

Weel et al. (2013) recently described an egg-based system extending the seed-based protocol from the previous section. The idea is that some of the robot modules that are not part of a robot organism act as an egg whose function is to collect and process genomes of robot organisms for reproduc-

tion. An egg is thus a stationary robot module that organisms can fertilize by sending their genome to it. An egg that has been fertilized by a number of organisms selects two of the received genomes for recombination followed by mutation. Then the egg becomes a seed, and initiates the morphogenesis of a new organism using the new genome.

This system has been implemented in a rather simple, fast simulator, RoboRobo³ and numerous experiments have been conducted to gain insights into the ‘inner logic’ of this system. In particular, three major parameters have been identified: *egg lifetime*, i.e., how long eggs listen for genomes, *seed lifetime*, i.e., how long a fertilized egg (a seed) is allowed to build the organism its genome encodes before aborting, and *organism lifetime*, i.e., how long a fully grown organism lives before it dies. These experiments have disclosed how these parameters interact, in particular regarding their influence on the size of the organism population, the stability of the organism population, and the average size of the organisms.

Mature Life: Self-Repair

There are many complex steps proposed in this paper: birth, infancy, mature life over a sustained period. All of these complex and potentially error prone steps may well cause, or be inhibited by faults. Hence, throughout the lifetime of the robotic system, it is inevitable that there will be some form of failure within a robot, or within the organism. When such failures occur, the ability of the organism to perform its task, or even survive, is compromised. Failures can be caused by a range of different faults ranging from mechanical failures, to electronic hardware or software faults and as such prevent the organism from performing its task. For continued operation over the full lifetime of the robot/organism some form of self-repair is needed. The ability to self-repair is linked to the ability of the organism to perform morphogenesis, as it is very likely that some form of reconfiguration is needed in the event of failure.

We report here on two approaches of self-repair that have been explored. The first could be considered a type of self-assembly, as reported in Murray et al. (2013) where robots are able to form ad-hoc structures, with no pre-determined shape, as opposed to work described above where a shape is seeded into the robotic unit. Murray et al presented an algorithm that showed successful reconfiguration ability of specifically tailored e-pucks that could form the aforementioned structures.

Further work by the SYMBRION project, as yet unpublished, goes much further to permit a true self-repair approach for organisms. Using techniques developed within the project for the detection Timmis et al. (2010) and diagnosis Bi et al. (2010) of faults, combined with the morphogenesis approach described here, SYMBRION organisms can

perform a partial disassembly then a full reassembly back to the original structure, in a distributed and autonomous manner. Should a robotic unit fail at any position within the organism, the approach permits for the removal of that unit and a reconstruction of the organism using the morphogenesis approach described.

Concluding Remarks

In this paper we have introduced the Triangle of Life: a conceptual framework for artificial systems in which robots actually reproduce. Our proposed framework contrasts with traditional evolutionary robots approaches in several ways. Firstly, the life cycle does not run from birth to death, but from conception (being conceived) to conception (conceiving one or more children). Secondly we envision the whole process taking place in real time, with real robots in the real world. We do not prescribe how the process should be implemented, but two contrasting approaches present themselves: one in which some infrastructure provides materials and processes for robot birth, and another infrastructure-less approach which could be thought of as an extension to modular self-assembling robotics. The third departure from conventional practice is that fitness is tested primarily through survival to maturity and successful mating, rather than against an explicit fitness function. Thus a large number of factors including individual health and development, the living environment (which may include multiple generations of conspecifics), and simple contingency will influence whether an individual survives to pass on its genetic material. Importantly it follows that selection is also implicit. Although we are describing an artificial life system, the process of selection is much closer to Darwinian natural selection.

Finally we should speculate on how such an artificial life system might be used. Two contrasting applications present themselves. One as an engineering solution to a requirement for multiple robots in extreme unknown or dynamic environments in which the robots cannot be specified beforehand: robots required to explore and mine asteroids, for instance. The other application is scientific. Our proposed artificial life system could be used to investigate novel evolutionary processes, not so much to model biological evolution – life as it is, but instead to study life as it could be.

Acknowledgements

The authors gratefully acknowledge the contributions of all SYMBRION partners. In particular, we are indebted to Markus Dauschan, Evert Haasdijk, Serge Kernbach, Wenguo Liu, Jean-Marc Montanier, Jessica Meyer, Lachlan Murray, Marc Schoenauer, Michele Sebag, and Berend Weel.

References

- Bentley, P. and Kumar, S. (1999). Three ways to grow designs: A comparison of embryogenies for an evolutionary design problem. In et al., W. B., editor, *GECCO'99*, pages 35–43. Morgan Kaufmann.

³<https://code.google.com/p/roborobo/>

- Bi, R., Timmis, J., and Tyrrell, A. (2010). The diagnostic dendritic cell algorithm for robotic systems. In *12th IEEE Congress on Evolutionary Computation (CEC10)*, pages 4280–4287.
- Bongard, J. and Pfeifer, R. (2003). Evolving complete agents using artificial ontogeny. In Hara, F. and Pfeifer, R., editors, *Morpho-functional Machines: The New Species (Designing Embodied Intelligence)*, pages 237–258. Springer-Verlag, Berlin.
- Dauschan, M., Thenius, R., Schmickl, T., and Crailsheim, K. (2011). Using virtual embryogenesis multi-robot organisms. In *International Conference on Adaptive and Intelligent Systems, ICAIS'11, Klagenfurt, AT, September 06-08, 2011. Proceedings*, pages 238–247.
- Devert, A., Bredeche, N., and Schoenauer, M. (2006). Blind-builder: a new encoding to evolve lego-like structures. In Collet, P., Tomassini, M., Ebner, M., Gustafson, S., and Ekart, A., editors, *Proceedings of EUROGP 2006*, pages 61–72. Springer.
- Devert, A., Bredeche, N., and Schoenauer, M. (2011). Robustness and the halting problem for multicellular artificial ontogeny. *IEEE Trans. Evolutionary Computation*, 15(3):387–404.
- Doursat, R., Sayama, H., and Michel, O., editors (2012). *Morphogenetic Engineering: Toward Programmable Complex Systems*. Springer series: Understanding Complex Systems. Springer.
- Eiben, A.E., Kernbach, S., and Haasdijk, E. (2012). Embodied artificial evolution – artificial evolutionary systems in the 21st century. *Evolutionary Intelligence*, 5(4):261–272.
- Funes, P. and Pollack, J. (1997). Computer evolution of buildable objects. In Husbands, P. and Harvey, I., editors, *Fourth European Conference on Artificial Life*, pages 358–367. MIT Press.
- Langdon, C., editor (1989). *Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems (ALIFE '87)*. Santa Fe Institute Studies in the Sciences of Complexity, Addison-Wesley.
- Langton, C., editor (1995). *Artificial Life: an Overview*. MIT Press, Cambridge, MA.
- Levi, P. and Kernbach, S., editors (2010). *Symbiotic Multi-Robot Organisms: Reliability, Adaptability, Evolution*, volume 7 of *Cognitive Systems Monographs*. Springer, Berlin, Heidelberg, New York.
- Levy, S. (1992). *Artificial Life*. Vintage Books.
- Lipson, H. and Pollack, J. B. (2000). Automatic design and manufacture of robotic lifeforms. *Nature*, 406:974–978.
- Liu, W. and Winfield, A. F. (2012). Distributed autonomous morphogenesis in a self-assembling robotic system. In Doursat, R., Sayama, H., and Michel, O., editors, *Morphogenetic Engineering*, Understanding Complex Systems, pages 89–113. Springer Berlin Heidelberg.
- Lund, H. (2003). Co-evolving control and morphology with lego robots. In *Morpho-functional Machines: The New Species*, pages 59–79. Springer.
- Murata, S. and Kurokawa, H. (2012). *Self-Organizing Robots*, volume 77 of *Springer Tracts in Advanced Robotics*. Springer.
- Murray, L., Timmis, J., and Tyrrell, A. (2013). Modular self-assembling and self-reconfiguring epucks. *Swarm Intelligence*. (in press).
- Nolfi, S. and Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press, Cambridge, MA.
- Peysakhov, M., Galinskaya, V., and Regli, W. C. (2000). Using graph grammars and genetic algorithms to represent and evolve lego assemblies. In *Proceedings of GECCO 2000*, pages 269–275.
- Stanley, K. and Miikkulainen, R. (2003). A taxonomy for artificial embryogeny. *Artificial Life*, 9(2):93–130.
- Stradner, J., Hamann, H., Zahadat, P., Schmickl, T., and Crailsheim, K. (2012). On-line, on-board evolution of reaction-diffusion control for self-adaptation. In Adami, C., Bryson, D. M., Ofria, C., and Pennock, R. T., editors, *Proceedings of the ALife13*, pages 597–598. MIT Press.
- Thenius, R., Bodi, M., Schmickl, T., and Crailsheim, K. (2010). Using virtual embryogenesis for structuring controllers. In Hart, E., McEwan, C., Timmis, J., and Hone, A., editors, *Artificial Immune Systems, 9th International Conference, ICARIS 2010*, volume 6209 of *Lecture Notes in Computer Science*, pages 312–313. Springer.
- Timmis, J., Tyrrell, A., Mokhtar, M., Ismail, A., Owens, N., and Bi, R. (2010). An artificial immune system for robot organisms. In Levi, P. and Kernbach, S., editors, *Symbiotic Multi-Robot Organisms: Reliability, Adaptability, Evolution*, pages 268–288. Springer.
- Trianni, V. (2008). *Evolutionary Swarm Robotics – Evolving Self-Organising Behaviours in Groups of Autonomous Robots*, volume 108 of *Studies in Computational Intelligence*. Springer.
- Waibel, M., Keller, L., and Floreano, D. (2009). Genetic Team Composition and Level of Selection in the Evolution of Cooperation. *IEEE Transactions on Evolutionary Computation*, 13(3):648–660.
- Wang, L., Tan, K., and Chew, C. (2006). *Evolutionary Robotics: from Algorithms to Implementations*, volume 28 of *World Scientific Series in Robotics and Intelligent Systems*. World Scientific.
- Weel, B., Haasdijk, E., and Eiben, A.E. (2013). Body building: Hatching robot organisms. In *IEEE International Conference on Evolvable Systems (IEEE ICES 2013)*. IEEE Press.
- Yim, M., Shen, W., Salemi, B., Rus, D., Moll, M., Lipson, H., Klavins, E., and Chirikjian, G. (2007). Modular self-reconfigurable robot systems. *IEEE Robotics and Automation Magazine*, 14(1):43–52.
- Zykov, V., Mytilinaios, E., Desnoyer, M., and Lipson, H. (2007). Evolved and designed self-reproducing modular robotics. *IEEE Transactions on Robotics*, 23(2):308–319.