18

FROM *R.U.R.* TO ROBOT EVOLUTION

Geoff Nitschke and Gusz Eiben

R.U.R: Rossumovi Univerzální Roboti (*Rossum's Universal Robots*) presents a narrative that discusses age-old postulates for what separates humankind and automata. For example, at the end of the play, Alquist says: "If you want to live, then mate like animals!" The idea of robots breeding like animals must have sounded ludicrous in 1920, and even in 2020 it sounds far-fetched, if not impossible. How will it sound in 2120? After all, the difference between us and our mechanical servants is that robots are *made, not born*.[1] Thus, the key question is: *Can robots have children?*

Self-replication has been a long-standing open research problem and topic of discussion in artificial life,[2] with a range of highly anticipated future macrorobotic to nanorobotic applications.[3] More recently, self-replication has been the subject of some research attention in the field of evolutionary robotics,[4] and the topic has even enjoyed some international media attention.[5]

However, to date evolutionary robot systems are almost all in simulation and primarily concerned with evolving the brains (i.e., controllers) of the robots, not with evolving the bodies (i.e., morphologies). Hence, the evolvable entities that reproduce and get selected are inherently digital, living in a virtual world. Even the most prominent papers are limited in this respect. For instance, in the system described by Hod Lipson and Jordan Pollack in 2020,[6] the evolution of robots takes place in computer simulation and only one robot, the best result of the evolutionary process, is produced in the real world. As of 2022, real robots do not seem to be able to reproduce and evolve. Evolutionary robotics is a multidisciplinary research field drawing from embodied artificial intelligence,[7] cognitive science, evolutionary biology, evolutionary computing, and robotics, and also has significant crossover with artificial life. An end goal of evolutionary robotics, and more generally of artificial life, is to apply biologically inspired principles as adaptive mechanisms in designing artificial organisms.[8] After some three decades of experimental evolutionary robotics research, many have become frustrated by the limitations of current behavioral adaptation approaches in physical robots.[9] Such limitations emerge from applying one's machine learning method of choice to adapt only the incorporeal controller program constructs (software encoding sensory-motor correlations) that instantiate robot behavior within static corporeal bodies (hardware comprising sensory-motor configurations).

Currently, adaptability in evolutionary robotics assumes the form of robot controllers learning behaviors suitable for solving specific tasks in specific environments. With few notable exceptions,[10] the physical chassis, sensors, actuators, motors, and power source defining the bodies (morphologies) of robots are fixed, and any morphological adaptations are implemented as time-consuming manual reconfiguration of sensory-motor systems by human engineers. Thus, adaptability in current experimental robots is significantly limited by their morphologies.[11] In contemporary evolutionary robotics, this means that robots designed for specific environments are only adaptable to tasks within those environments.

In the fictional world of *R.U.R.*, Čapek's robots have replicas of human bodies[12] and, importantly, this morphology gives them a generalist (universal) capability to operate in many types of environments across the world and to perform a vast range of tasks. Even though the *R.U.R.* robots are behaviorally wired for specialized assigned tasks, their generalist human bodies give them the capability to potentially learn any skill or behavior or accomplish any task, just as their human creators could. However, a point of contention for the *R.U.R.* robots is that unlike their human creators, they are unable to self-propagate and improve themselves over the evolutionary timescale of generations.

In *R.U.R.* this was an allusion to the assembly-line nature of the robots, where each robot was simply a product, manufactured for a specific purpose. This is akin to contemporary industrial robots used in global manufacturing industries, performing repetitive tasks for their lifetimes, until their physical components degrade, or their power source is exhausted, resulting in disposal by their manufacturers. Unlike their human creators, the *R.U.R.* robots did not contain a blueprint of themselves (a robotic DNA), meaning reproduction was impossible, which relegated the robots to the ranks of biomechanical chattel rather than artificial life.

To the best of our knowledge the earliest paper that specifically addressed the artificial evolution of physical entities (artificial life forms) is that published by Gusz Eiben and his colleagues in 2012.[13] This paper discusses the *evolution of things* and presents a list of four properties that distinguish a strongly embodied evolutionary system from mainstream evolutionary computing and evolutionary robotics. First, a strongly embodied evolutionary system uses physical units instead of just virtual individuals. Second, there is real birth and death, where reproduction creates new (physical) objects, and survivor selection eliminates some of them. Third, artificial evolution is driven by environmental selection or a combination of environmental fitness and user-defined task-based fitness. Fourth, reproduction and survivor selection are not coupled by an overseeing "manager" as usual in evolutionary computing and simulation-based evolutionary robotics, but are executed in a distributed manner (individuals can decide themselves who to reproduce with). Reproduction of physical artifacts is one of the grand challenges this paper identifies.

This challenge has been treated in a proposed generic system architecture: the Triangle of Life.[14] The Triangle of Life constitutes a robotic life cycle that runs from conception (being conceived) to conception (conceiving offspring). This triangle consists of three stages: morphogenesis, infancy, and mature life. The second stage, infancy, is an important new element in a generic robot evolution framework. In this stage newborn robots undergo (supervised) learning to acquire and optimize essential skills required in the given environment and for the tasks at hand. Robots that fail to learn the necessary skills are removed from the system to prevent reproduction of inferior robots and save resources. Robots that successfully learn these skills become fertile adults and can reproduce. To this end, the mate selection mechanism can be innate in the robots, but depending on the application it can also be executed by an overseer, which can be algorithmic or a human breeder.

The problem of robot birth is handled in the morphogenesis stage of the triangle. The main idea behind robot reproduction is to follow nature's solution and distinguish two levels of existence in each robot: the genotype and the phenotype. The genotype is the code (DNA in nature, or a specification sheet in robotics), whereas the phenotype is the physical expression of this code (the real animal or robot). By this distinction robot reproduction can be decomposed into two steps: (1) introducing variation in the genotypes and (2) constructing a new robot phenotype encoded by a given genotype. It is important to note that the genotypes are digital entities, pieces of computer code that can be manipulated easily. In particular, we can use mutation of crossover operators from evolutionary computing[15] such that a new genotype, and hence corresponding robot phenotype, inherits its parents' characteristics. As for the second step, the actual birth can be instantiated via employing 3D printing, automated assembly, or both to construct the phenotype encoded by a given genotype.

Recently, the first large-scale robot evolution project commenced. The Autonomous Robot Evolution (ARE) project is a collaboration between four universities in the UK and the Netherlands that aims to develop a fully operational EvoSphere,[16] that is, a robot system that implements all components of the Triangle of Life.[17] A key innovation of ARE is the deep integration of virtual and physical robot evolution into a hybrid evolutionary system.[18] Specifically, two concurrently running implementations of the Triangle of Life are envisaged, one in a virtual environment and one in the physical world, where the robot population evolving in the real world is assisted by a virtual population for efficiency. The essential feature behind the integrated system is the use of the same genetic representation in both virtual and physical worlds. This allows cross-fertilization (mating between virtual and physical robots) and twin creation (sending a robot's genotype to the other world to be constructed). Such an integration of the virtual and the physical subsystems offers the best of both worlds. Physical evolution is accelerated by the virtual component finding useful robot subsystems using less time and resources, while simulated evolution is accelerated by favorably tested physical robotic subsystems.

The idea of developing robot systems that reproduce and evolve in real time and real space has a twofold motivation. First, such systems are interesting from an engineering perspective. Evolution can be employed as a design method for complex environments and tasks where adequate morphologies and controllers cannot be obtained by traditional approaches. Evolving robots in the real world, not in simulation, is important to avoid the inevitable reality gap, the effect that the simulated and the real-world behavior of the evolved robots are very different.[19] Robots can then be developed through iterated selection and reproduction cycles until they satisfy the users' criteria. This technology amounts to robot breeding: the user drives evolution and can stop when an optimal solution is found. Further to optimizing designs, evolution has the ability to adapt to unknown and changing conditions, for example, in space research or the exploration of remote areas on Earth. An evolving robot population can adapt to the given circumstances and repeatedly readjust if the conditions change. Overall, an evolution-of-things technology will allow for radically new types of machines, able to adapt their form and function, possibly without direct human oversight.

Second, evolutionary robot systems provide a new approach for scientific research. Akin to a telescope used in astronomy research or a cyclotron needed to study nuclear particles, an EvoSphere where robots reproduce and evolve forms a novel research instrument for studying evolution. Using robots offers important advantages with respect to biological experimentation: the experimental conditions are easy to control, robot characteristics can be observed and logged easily, system properties can be simply fixed, and several repetitions can be done for statistical purposes. It can be argued that artificial evolution can and will differ from natural evolution. However, this need not be a problem, rather an opportunity. This vision has been eloquently phrased by the evolutionary biologist John Maynard Smith,[20] and discussed in a grand perspective of natural and artificial evolution: "So far, we have been able to study only one evolving system and we cannot wait for interstellar flight to provide us with a second. If we want to discover generalizations about evolving systems, we will have to look at artificial ones."[21]

Robots that reproduce and evolve *in the wild* can represent a danger, as speculated in science fiction.[22] Specifically, a runaway evolution scenario, where uncontrolled and unlimited reproduction leads to large numbers of potentially dangerous robots, should be prevented. For reasons of ethics and safety such issues must be considered from the very beginning of the development. Specifically, one should not set up a physically evolving robot system without an emergency switch, that is, a fail-safe way of stopping the system. One particular solution can be the denial of all distributed reproduction systems, for instance, the robotic equivalents of cell division, laying eggs, or pregnancy. Instead we should use a centralized and externalized reproduction mechanism, a distinguished infrastructure, such as a robotic birth clinic or production center. From the robots' point of view this is a single point of failure; for us humans this is the *kill switch*. If we shut down the robot reproduction center, we effectively shut down evolution. The Triangle of Life architecture and EvoSphere concept are naturally suited for this, and while in principle there can be other approaches to guaranteeing safety, we recommend using centralized external reproduction centers and being wary of distributed alternatives.

As for some concluding remarks on the topic, let us recall the question from the introduction: *Can robots have children?* The answer is positively yes. The field is in an embryonic stage, but there is a large body of knowledge regarding digital evolution, and crossing the border to physically embodied evolution seems imminent. The main obstacle at the moment is the 3D printing and rapid prototyping technology. As of today it is not possible to print a fully functional robot, except very simple ones, but the technology is developing quickly. The ability to print motors, CPUs, wires, sensors, and various actuators resulting in the automated production and assembly of robots is hypothesized to become reality in the coming decade. In the fictional world of *R.U.R.*, manufactured robots could essentially replace many of the current societal roles assumed by human workers. In evolutionary robotics, artificially evolved robots would ideally be general enough to adapt to a vast range of environments, or be plastic enough to morphologically and behaviorally adapt over successive generations as robots move between environments. More importantly, such robots should have the capability to self-replicate as well as evolve and learn, giving them the broadest possible spectrum of adaptability. To achieve such flexibility and generality in robot body-brain artificial evolution, these fundamental biological mechanisms must be formulated as a methodology for automated robot design. As in *R.U.R.* where robots are directed to manufacture other robots, such an automation methodology would ideally be embodied as an automated robot-design factory, situated in any environment for the purpose of evolving robots optimally suited to solving tasks within that environment.

Automated robot-design factories (founded on the Triangle of Life architecture and the EvoSphere concept) would need to be driven by suitable general user-specified goals akin to current scientific and industrial mission directorates such as discovering traces of water on other planetary bodies or potential oil and gas mining sites. Rapid prototyping and 3D printing technologies[23] could then be applied in the context of the factory iteratively, producing generations of physical robot prototypes. Each robot in each generation would act in its environment, be evaluated by the factory, and receive a fitness score proportional to its task performance. Subsequently, the current generation would then be decommissioned and recycled for materials and components. The fittest robots of this current generation would then have their controllers, materials, and components reused and recombined to produce the next generation of improved robotic (body-brain) designs. A cycle of robot production, exploration, information gathering, body-brain redesign, decommissioning of robots for recycling, reuse and recombination of their parts in the next stage of robot production could continue indefinitely. Practically, this evolutionary design, production, and evaluation cycle would continue for the duration of factory power sources or until the robots perfectly adapt to their environment and tasks. Such an artificial life counterpart to nature, AutoFac, was recently proposed,[24] in which generations of robots emerging from such robot factories would embody evolving controller-morphology designs. Importantly, such artificial evolution runs at orders of magnitude faster than natural evolution, expedited by reuse of engineered designs.

Finally, this raises the question of why exactly, and for what types of task environments, we would need such an automated robot-design factory. The key envisaged benefit is that such a factory would be a mobile design and production center that could be dropped, as a problem-solving tool, into any remote and hostile environment.[25] The factory could then automatically produce robot colonies as solutions in response to complex and arduous tasks for which we have little to no a priori knowledge. Specifically, it could support tasks in unpredictable, dynamic, and unknown environments where optimal robot body-brain architecture cannot be preengineered due to complexities in the environment, and where it is beneficial for robots to adapt their own morphologies and behaviors as they explore the environment.

Example applications include space exploration,[26] search and rescue,[27] disaster management,[28] environmental monitoring[29] and asteroid mining.[30] In the world of *R.U.R.*, the robots were similarly produced as problem-solving tools. The true potential of fully automated, self-propagating and self-adapting robotic systems will be in unexplored or remote, inhospitable environments, too hazardous for humans to live and work in, accomplishing tasks we ourselves do not know how to solve. As a complement to the capitalistic objective of the *R.U.R.* robot creators to reduce manufacturing costs and increase profits, consider the enormity of potential scientific discoveries and industrial gains to be gotten from deploying fully automated robot-design factories as solutions to unsolved tasks across a plethora of environments. We anticipate that fully automated, self-sustaining and artificially evolving robot colonies will become indispensable problem-solving tools enabling us to solve increasingly complex problems and problems that we currently cannot imagine.