

This is a repository copy of *The ARE Robot Fabricator:How to (Re)produce Robots that Can Evolve in the Real World*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/154551/>

Version: Accepted Version

---

### **Proceedings Paper:**

Hale, M, Buchanan Berumen, Edgar orcid.org/0000-0001-6587-8808, Winfield, Alan et al. (5 more authors) (2019) *The ARE Robot Fabricator:How to (Re)produce Robots that Can Evolve in the Real World*. In: *International Society for Artificial Life:ALIFE2019*. , pp. 95-102.

---

### **Reuse**

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

### **Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.

# The ARE Robot Fabricator: How to (Re)produce Robots that Can Evolve in the Real World

Matthew F. Hale<sup>1\*</sup>, Edgar Buchanan<sup>2\*</sup>, Alan F. Winfield<sup>1</sup>, Jon Timmis<sup>2</sup>, Emma Hart<sup>4</sup>, Agoston E. Eiben<sup>3</sup>, Mike Angus<sup>2</sup>, Frank Veenstra<sup>4</sup>, Wei Li<sup>2</sup>, Robert Woolley<sup>2</sup>, Matteo De Carlo<sup>3</sup> and Andy M. Tyrrell<sup>2</sup>

<sup>1</sup>Bristol Robotics Laboratory, University of the West of England, Bristol, UK

<sup>2</sup>Department of Electronic Engineering, University of York, UK

<sup>3</sup>Department of Computer Science, Vrije Universiteit Amsterdam, NL

<sup>4</sup>School of Computing, Edinburgh Napier University, UK

matt.hale@brl.ac.uk, edgar.buchanan@york.ac.uk \*These authors contributed equally to this paper

## Abstract

The long term vision of the Autonomous Robot Evolution (ARE) project is to create an ecosystem of both virtual and physical robots with evolving brains and bodies. One of the major challenges for such a vision is the need to construct many unique individuals without prior knowledge of what designs evolution will produce. To this end, an autonomous robot fabrication system for evolutionary robotics, the *Robot Fabricator*, is introduced in this paper. Evolutionary algorithms can create robot designs without direct human interaction; the Robot Fabricator will extend this to create physical copies of these designs (phenotypes) without direct human interaction. The Robot Fabricator will receive genomes and produce populations of physical individuals that can then be evaluated, allowing this to form part of the evolutionary loop, so robotic evolution is not confined to simulation and the reality gap is minimised. In order to allow the production of robot bodies with the widest variety of shapes and functional parts, individuals will be produced through 3D printing, with prefabricated actuators and sensors autonomously attached in the positions determined by evolution. This paper presents details of the proposed physical system, including a proof-of-concept demonstrator, and discusses the importance of considering the physical manufacture for evolutionary robotics.

## Introduction

This paper outlines a long-term vision towards robots that reproduce and evolve in real-time and real space as well as an ongoing research project concerned with the first tangible implementation of such robots; evolving physical robots will be a significant step towards robotic artificial life. Specifically, we discuss the challenge of robot (re)production and present our first results with the ‘Robot Fabricator’ for autonomously producing robot phenotypes.

The long-term vision behind this research has a two-fold motivation. Firstly, it is to create a new type of artificial evolutionary systems that depart from the evolution of digital artefacts—Evolutionary Computing—and realizes the evolution of physical artefacts: the Evolution of Things as introduced in (Eiben et al., 2012). Such systems will represent a third incarnation of Darwinian principles. To date, these principles can be observed and studied in wetware (Life on

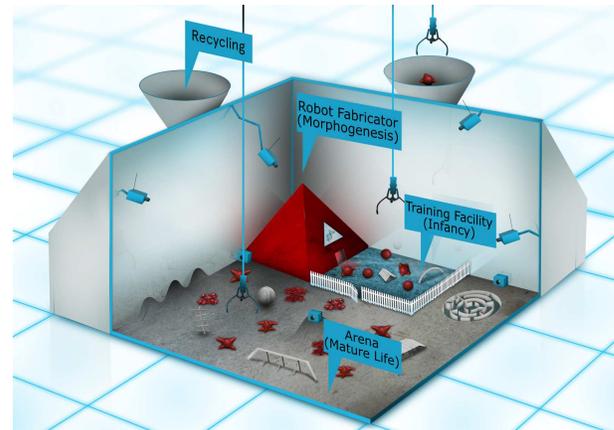


Figure 1: Illustration of the ARE environment, showing the three main stages of the Triangle of Life model and a recycling facility.

Earth) and software (Evolutionary Computing); the Evolution of Things will realize them in hardware, cf. Eiben and Smith (2015). Such hardware models of evolution will facilitate fundamental research into, for instance, the macro-level mechanisms of evolution, the emergence of (embodied) intelligence, and the simultaneous evolution of the body and the brain without suffering from the infamous reality gap (Jakobi et al., 1995).

Secondly, evolving robots is interesting from an engineering perspective. The fact that Life on Earth has populated practically all possible environmental niches demonstrates that natural evolution is very successful in producing specialised life forms. Hence, it is a reasonable hypothesis that artificial evolution will be capable of producing specialised robots for various environments and tasks. Furthermore, an autonomously evolving robot population has the ability to adapt to previously unknown and/or changing conditions, thus creating new types of machines that are able to adapt their form and behavior.

The field of Evolutionary Robotics has addressed the evolution of robot controllers (brains) with considerable suc-

cess but evolving the morphologies (bodies) has received much less attention (Nolfi et al., 2000, 2016; Doncieux et al., 2015). This is somewhat understandable, given the difficulty of implementing such systems in the real world, i.e. the lack of technologies for automated (re)production of robots. However, advances in robotics, 3D-printing, and automated assembly mean it is now timely to work on physical robot systems with evolvable morphologies (Winfield and Timmis, 2015). The Autonomous Robot Evolution (ARE) project<sup>1</sup> is concerned with developing the first such system, illustrated as a concept in Figure 1. The work will allow for radically new autonomous systems, where robots are designed and manufactured by algorithms and machines rather than by humans.

The main contribution of this paper is the proposal of a novel approach to robotic evolution, where an automated facility enables the (re)production of robot populations, and thus robot evolution in the physical world. After discussing previous related work on evolving physical robot morphologies, the paper outlines the overall system architecture for the ARE project. This is followed by a more detailed description of the required physical infrastructure, focusing on the proposed Robot Fabricator system (RoboFab) to achieve an automated assembly process, with a proof-of-concept demonstration presented.

## Related Work

As of today there have been very few examples of autonomous robot fabrication for evolutionary robotics in the literature. Most approaches involve evolving the robots in different simulators and manually assembling the resulting robots. In addition, those robots that are autonomously assembled have generally been very simple with a single type of actuator and no sensors.

The important link between morphology and control (Pfeifer and Bongard, 2007) suggests that they should be evolved together in order to unlock the potential of evolutionary robotics; however since the breakthrough work of Sims (1994), the majority of brain and body evolution has occurred in simulation.

A logical progression from simulation studies has been the simulate-then-transfer paradigm, such as the Golem project (Pollack and Lipson, 2000; Lipson and Pollack, 2000), in which the evolution of simple robots (without sensors and with fixed controllers) occurred in simulation after which a select few individuals were physically manufactured. However this does not allow for any selection based on the physical robots, and so evolution stops at the point of physical manufacture. Although this paradigm has been shown to be successful for problems such as antenna design (Hornby et al., 2011), the complex interactions of mobile robots with their environment are difficult to simulate

accurately, leading to the reality gap (Jakobi et al., 1995), which could be bypassed by evaluating robots in the physical world.

Evolution of controllers where virtual models of the robots are updated according the performance of physical robots has been achieved previously, e.g. (Bongard and Lipson, 2004; Hwangbo et al., 2019).

One approach for creating populations of physical robots may be through self-replicating robots. Zykov et al. (2005, 2007) demonstrated modular robots capable of such reproduction. However the physical robots only created exact copies of themselves, without variation and therefore without evolution in hardware.

There has been some progress toward a fully automated system for creating physical robot phenotypes (which in ARE is termed the Robot Fabricator). Brodbeck et al. (2015) used a “mother robot” (a robot arm) to create physical individuals by gluing cuboid modules together. They demonstrated model-free morphological evolution, with 100 individuals for each evolutionary run physically created and tested; a major breakthrough. However, the robots used were limited to a single type of actuator (servo-motor) and the robots had no sensors—evolving more complex morphologies is likely to present further challenges.

3D printing is a rapidly developing field, offering exciting new opportunities to create complex robotics systems with minimal manual intervention (MacCurdy et al., 2016), with a particular interest in printed soft robots (Hiller and Lipson, 2012; Bartlett et al., 2015). However, 3D printing cannot yet be used to create mechatronic components, making essential robot functions of sensing and actuation impossible. This motivates a combined approach in the ARE project, where some components are 3D printed (allowing arbitrary shapes), while sensor and actuator ‘organs’, with electronics and motors, are hand designed and built.

The Triangle of Life concept (discussed below) has been demonstrated with a simplified setup by Jelisivcic et al. (2017), where physical robots were 3D-printed and hand-assembled, also using prefabricated organ, such as servo motors and Raspberry Pis. The ARE project will take this further through the automated Robot Fabricator, allowing larger numbers of individuals and a complete evolutionary system to be fabricated.

## Overall System Architecture

A general architecture for evolving robots in real time and real space has been suggested in the conceptual framework named the Triangle of Life by Eiben et al. (2013), shown in Figure 2. A real-world implementation of this is envisaged by the notion of an EvoSphere as introduced and extensively discussed in Eiben (2015) and modified for the ARE project in Figure 1; this forms a design template for an evolutionary robot habitat and provides the basis of the physical environment in the ARE project.

<sup>1</sup>see <https://www.york.ac.uk/robot-lab/are/>

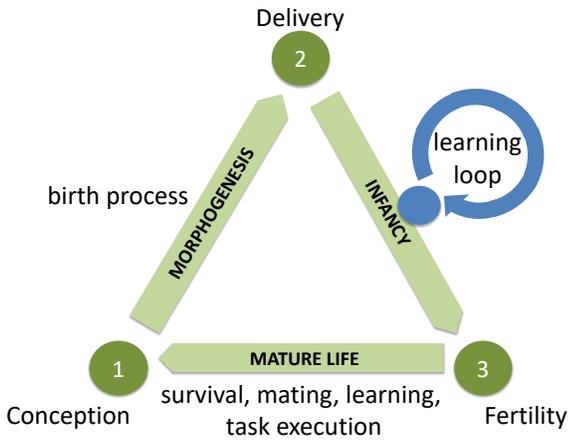


Figure 2: Generic system architecture for robot evolution conceptualized by the Triangle of Life. The learning methods in the Infancy stage are not necessarily evolutionary.

The Triangle of Life consists of three stages: morphogenesis, infancy, and mature life, as illustrated in Figure 2. Consequently, an EvoSphere consists of three components. The Robot Fabricator is where new robots are created (morphogenesis). The Training Facility provides a suitable environment for individuals to learn during infancy, providing feedback, perhaps via a computer vision system and/or a human user, so individual robots can learn to control their (possibly unique) body to perform some simple tasks. The Training Facility increases the chances of success in the Arena and plays an important role: it prevents reproduction of poorly performing robots and saves resources. It also enables refinement of controllers learned in a simulated environment that do not transfer properly due to the reality gap. If a robot acquires the required set of skills, it is declared a fertile adult and enters the Arena, which represents the world where the robots must survive and perform user-defined tasks, and may be selected for reproduction. The selection mechanism can be innate in the robots (by choosing “mates”) or executed by an overseer, which can be algorithmic or a human “breeder”.

An essential feature of the EvoSphere concept and the ARE system is the centralised, externalised reproduction. For reasons of ethics and safety we reject distributed reproduction systems, e.g. self-replicators or the robotic equivalents of cell division, eggs, or pregnancy, and deliberately choose for one single facility that can produce new robots. This facility, the Robot Fabricator, serves as an emergency switch that can stop evolution if the users deem it necessary.

The ARE system features deep integration of virtual and physical robot evolution. In essence, there are two concurrently running implementations of the Triangle of Life, one in a virtual environment and one in the physical environment. The Ecosystem Manager is a program to control the hybrid physical-virtual system, providing the link between

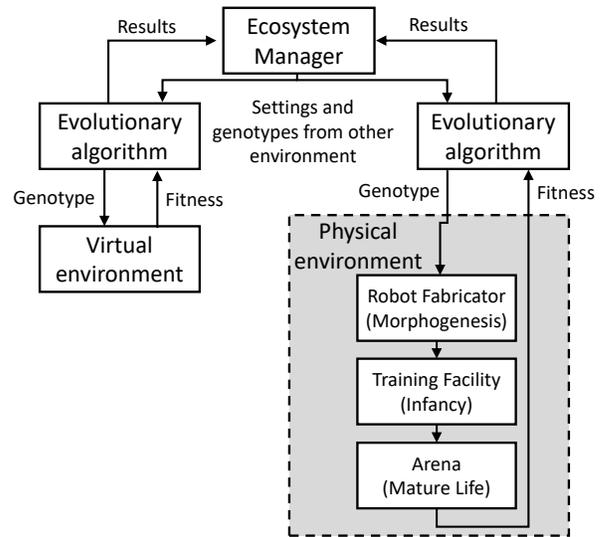


Figure 3: Diagram of the ARE system components: the virtual environment, the physical environment, and the Ecosystem Manager. Simulated and physical evolution are running concurrently, supervised and controlled by the Ecosystem Manager.

the two, as illustrated in Figure 3. This integration is made possible by using the same genetic representation in both worlds, enabling cross-breeding so a new robot in either environment could have physical or virtual parents, or a combination of both. An individual can also be copied between environments simply by transferring its genotype from one environment to the other. This integration must be controlled by the Ecosystem Manager, which optimises the working of the hybrid evolutionary system, maximising task performance. It reacts to flows of information from both subsystems, and to human-specified goals, either hand-directed or running autonomously.

This integration seeks to combine the advantages of the real and virtual worlds. Physical evolution is accelerated by the virtual component that can find good robot features with less time and resources than physical evaluation, while simulated evolution benefits from the influx of genes that are tested favourably in the real world. This means while a single physical robot is evaluated in the physical evolution, hundreds of simulated generations will be evaluated at the same time. As consequence of this a huge proportion of the population will be virtual. In order to compensate for this we will give greater weighting to results obtained from physical robots.

It is important to mention that the size of the population in the physical world will dynamically change over time, as new robots are added (from the virtual world or randomly created) and due to physical robots malfunctioning or being removed due to poor performance.

The evolutionary algorithms are kept separate, as shown in Figure 3, to allow for different configurations (of population size, mutation rate etc.) suited to these different environments, with settings dictated by the Ecosystem Manager.

### Autonomous Manufacturing in the Physical Environment

This section describes the proposed method for the physical manufacture of robots in the context of the ARE project.

The *Robot Fabricator* is a system to automate the process of morphogenesis: the conversion of genotypes into physical robot phenotypes, as outlined in Figure 3. The final objective is to remove human intervention from this process as far as possible, creating a manufacturing system for the autonomous production of complete robots, which will allow the evaluation of physical individuals to form part of the evolutionary process.

The proposed design for the Robot Fabricator is described below.

The *organs* are defined as active components which individual robots can use to perform their task(s). It is expected each robot will have a single “brain” organ, which contains the electronic control hardware and a battery to provide the robot with power. Other organs provide sensing and actuation, with their type and location being specified in the genome written by the evolutionary algorithm. These other organs will typically each comprise a single sensor or actuator, such that evolution is free to select the quantity and arrangement of the active parts of the robot. In practice, once a robot is no longer needed, the physical organs can be removed and recycled to create new individuals. Examples of physical and virtual organs are shown in Figures 4 and 5.

The *skeleton* is the structure that physically connects the organs together. Specifically, the skeleton will be comprised of thermoplastic and made by additive manufacture (3D printing). The skeleton is generally expected to be non-functional except for serving as a scaffold to hold the organs in place relative to each other.

In this way, the robots created will feature organs which are designed to be re-usable for many different designs and the skeleton, which will be made specifically for a particular individual. This organ and skeleton approach is flexible enough to allow for a wide range of sensors and locomotion methods, including wheel organs and/or joint organs for constructing robots with limbs. Furthermore, certain organs could also be designed such that some part(s) can be individually 3D printed and attached autonomously at the time of assembling the robot, for instance the radius (or even the shape) of wheels, or the fingers of a gripper, so that these aspects of the organs can also be evolved. As such, the Robot Fabricator is intended to allow for experiments where the search space can be large and diverse enough for fundamental research into evolutionary robotics (and perhaps evolution in general), and the robots capable enough to perform

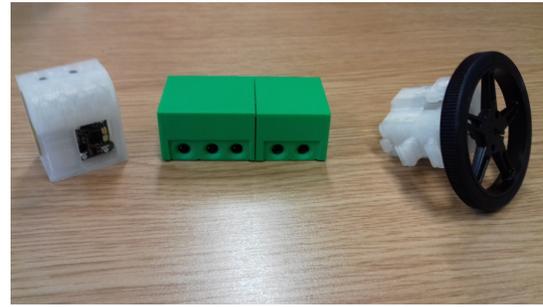


Figure 4: Prototype organs in the physical world. From left to right: a sensor organ, brain organ and wheel organ.

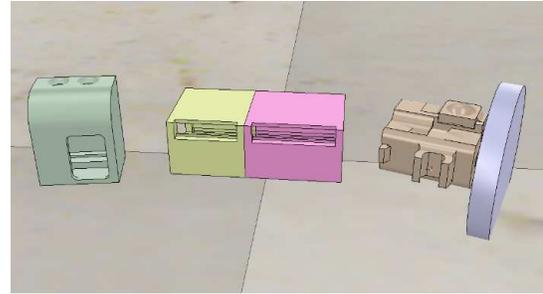


Figure 5: Models of the organs from Figure 4 for use in the virtual environment.

useful and interesting tasks.

One potential ethical issue is the waste produced by large numbers of bespoke plastic skeleton parts, which cannot be re-used. To minimise the environmental impact of this, the material chosen is a plant-based polymer (PLA), which is recyclable and biodegradable. After the useful life of an individual, the organs will certainly be removed and re-used for future individuals; it may also be possible to melt down the skeleton parts and make new filament, to be used for new individuals.

### The Robot Fabricator System

The skeleton, which holds the organs in their positions, needs to take an arbitrary shape depending on the organ locations specified. To achieve this, it is made using a 3D printer (LulzBot TAZ 6) by the Fused Deposition Modelling (FDM) approach to form 3D shapes from a thermoplastic (polylactic acid, commonly known as PLA).

The organs, in comparison, are much more complex, as they require electrical and electronic components and many different materials to allow for a range of actuation and sensing technologies; these will be prefabricated and attached to the skeleton using a multi-axis manipulator (a robot arm, in this case a Universal Robots UR5e).

To allow the robot arm to easily pick up each required organ, they are stored in an *organ bank*, where each organ is held in a known and accessible position.

While the organs are attached, the semi-constructed robot

must be held securely, but then must be released once the assembly is complete. This will be achieved by a fixture in the *assembly area* which can mate to a feature on the bottom face of the brain organ which forms the core of each individual.

The 3D printing of the skeleton is likely to take several hours, much longer than the assembly, and so the production rate can be increased by having multiple printers operating in parallel, tended to by a single robot arm. A proposed layout for these components is shown in Figure 6.

To demonstrate some of the key steps in this process, a test setup has been created. This process is depicted in Figure 7 and the supplementary video<sup>2</sup>. The manufacturing sequence is described as follows:

1. The Robot Fabricator receives the required coordinates of the organs and one or more mesh file(s) of the shape of the skeleton.
2. The skeleton is produced by the 3D printer (Figure 7-1).
3. The robot arm transfers the skeleton to the assembly area. This is currently done manually (Figure 7-2).
4. The organs are attached to the appropriate locations by the robot arm (Figure 7-3).
5. The robot arm connects the organs together by cables to provide power and communications (Figure 7-4).
6. Any remaining organs that cannot be attached by the robot arm, such as wheels, are attached manually (Figure 7-5). In the final design this step will be eliminated.
7. The robot is complete and ready to be transferred to the Training Facility (Figure 7-6).

## From a Virtual Robot to a Physical Robot

The robot morphologies shown in this paper were initially evolved in the simulations using a steady state evolutionary algorithm (EA), i.e. with overlapping generations. Traveling distance is the measure of fitness. The parameters values used for the EA are the following: the population size of 20, generation number of 200 and mutation rate of 0.2. An example of equivalent virtual and physical phenotypes is shown in Figure 8.

The robots shown in this paper were generated with direct encoding. The position and orientation of each organ is explicitly specified in the genome. The organs are connected directly to the brain with vertical and horizontal segments. This representation is used because of its simplicity, but further work will evaluate and compare other methods. The code that generates such morphologies can be found in the supplementary material of this paper.

Once the best individual is found, the list of coordinates of each organ is sent to the Robot Fabricator together with the mesh file that represents the shape of the body. The mesh file is used to create the skeleton with a 3D printer and the

list of coordinates is used to attach the organs to the body using the robot arm.

Even though a large diversity of robots with different shapes can be generated with evolutionary algorithms, not all the robots can be built. The practical limitations of the Robot Fabricator impose various constraints on what can be physically produced. Therefore, before a genome is sent to the Robot Fabricator, it is very important to make sure that its phenotype can be manufactured. This evaluation is described below.

## Viability Test

Before a genome is sent to the Robot Fabricator, it is subject to a *viability test* to make sure that its phenotype can be manufactured. Only viable robots will be produced, to avoid wasting time and resources attempting to manufacture only to find it cannot be made. Trying to produce a non-viable individual may also risk damage being caused to the organs or Robot Fabricator.

The viability test must check for violation of the limitations of the Robot Fabricator manufacturing process, which are constraints that would not exist for simulation, and would change depending on the assembly process chosen, such as those in Table 1.

The viability test may optionally be extended to cover not only individuals that *cannot* be manufactured, but also detect some cases of individuals with no chance of a decent performance, and therefore *should* not be manufactured. This could avoid wasting resources and speed up the overall evolution. For example, this could be a test of fundamental functionality—if a robot must sense its surroundings and act upon the information received, then it must possess at least one sensor and one actuator. Whether this type of addition to the viability test does significantly speed up evolution, or if they have negative side-effects, is an area to be explored.

In this paper we present some examples of robots generated with and without viability test from evolution. Robots are subjected to a series of checks when the the viability test is enabled, shown in Table 1; if a robot fails one of the checks it is considered as non-viable and a fitness of zero is assigned to this robot. All the checks are ignored and all the robots are considered viable when the viability test is disabled. An example of a robot that passes the viability test is shown in Figure 9.

Without the restrictions of the viability test, evolution will exploit any means of increasing the fitness function, which may not result in robots that can be manufactured. Figure 10 shows some examples of robots that evolved which would have failed the test, selected by manually inspecting the final generation to demonstrate different issues. Figure 10(a) shows an example of overlapping organs (although in this case a high fitness is achieved through a simulation bug, and is not the intended target of the viability test). With no limits to how many organs a robot can have, evolution can add

<sup>2</sup>Supplementary material available at: <https://www.york.ac.uk/robot-lab/are/alife2019/>

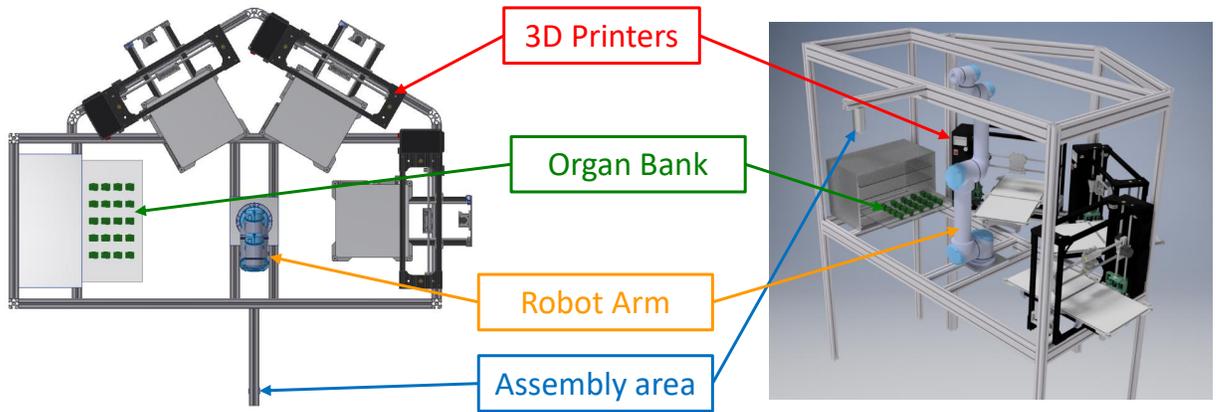


Figure 6: Concept layout for the final Robot Fabricator design to facilitate automated production of complete robots. The skeleton is printed on one or more printers, and assembled by the robot arm together with organs from the organ bank.

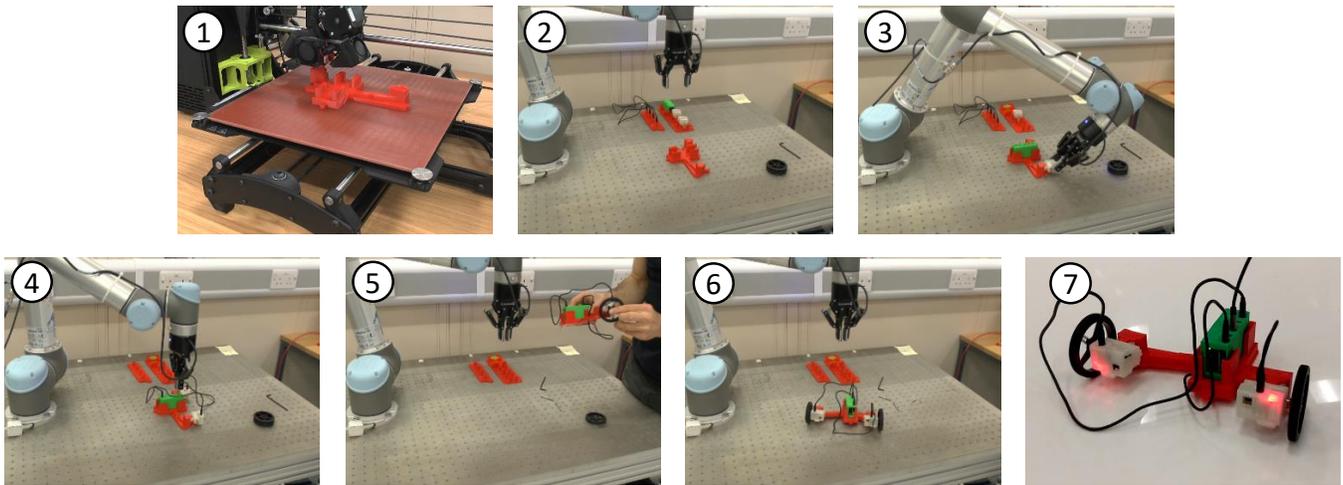


Figure 7: Robot production: (1) skeleton is 3D printed, (2), manual transfer to assembly area, (3) organs attached, (4) cables connected between organs, (5) manual wheel attachment, (6) robot is finished and (7) is tested. These images are from a video which is available online at [www.york.ac.uk/robot-lab/are/alife2019/](http://www.york.ac.uk/robot-lab/are/alife2019/).

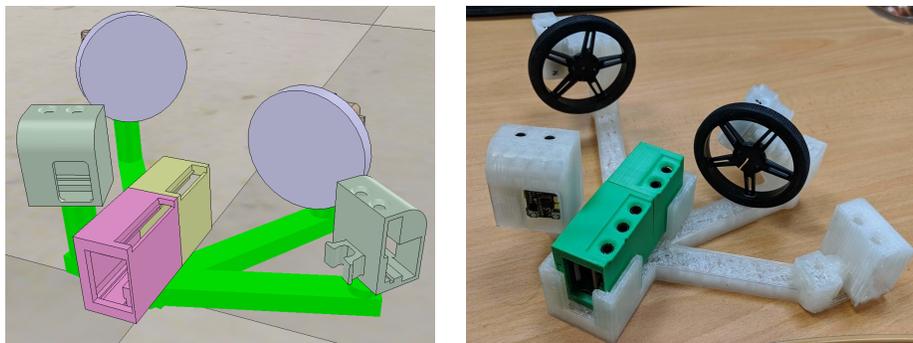


Figure 8: Equivalent virtual and physical phenotypes.

Constraint	Description
3D printer volume	The maximum size of a 3D-printed part ( <i>i.e.</i> a piece of skeleton) is limited by the print volume of the 3D printer. This does not prohibit several parts ( <i>e.g.</i> limbs) being made separately and then joined to form a complete robot larger than this.
Assembly area size	The overall size of any robot, after all the organs and limbs are connected, will be limited by the physical size of the assembly area.
Overlapping organs	In simulation organs can overlap and share space. In the physical world, this cannot be achieved.
3D Printing overhangs	3D printing works by building in layers; each layer must be supported by the layer below so there is a limit to the angle by which a face can overhang; additional support material would be difficult to remove autonomously.
Organ attachment and connections	The robotic arm requires access to attach each organ, so there must not be any 3D printed skeleton material, or a previously attached organ, in its path. Equivalent access is also required for each cable connection.
Number of Organs	The organs must be pre-made, and therefore there will be a fixed number available, limiting the quantity of organs that can be assigned to each individual.

Table 1: Manufacturing constraints to be considered by the viability test.

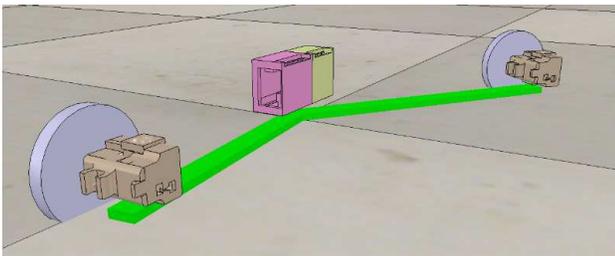


Figure 9: Robot evolved with viability test. This robot meets the physical constraints.

more than are available to the Robot Fabricator, such as the robot shown in Figure 10(b). In a similar way, evolution can exploit the size of the robot, resulting in skeleton sections larger than the print volume of the Robot Fabricator, as with the robot shown in Figure 10(c). Of course, some of these issues could be addressed by manipulating the fitness function or constraining the range of parameters within the genome, but it remains clear that there will be robots that are viable in a simulator but that cannot be manufactured for multiple reasons, therefore the viability test is essential.

It is worth mentioning that assigning a fitness of zero to a robot failing the viability test is not the only treatment for non-viable robots. Each robot could be repaired, or each robot could be kept in the population with low fitness, as it may be beneficial to allow movement through infeasible regions of the search space to find regions of higher fitness.

## Conclusions

The ARE project envisions an environment where autonomous systems (robots) are not designed by humans (or indeed designed at all), but are created through a series of steps that follow evolutionary processes. These robots will

be “born” through the use of 3D additive manufacturing, with novel materials and a hybridised physical-virtual evolutionary architecture. Newly created robots will learn in a safe and controlled environment where success will be rewarded. The most successful individuals will make available their genetic code for reproduction and for the improvement of future generations. Such a process may ultimately lead to a change in the way things are designed and manufactured.

This paper describes the first step towards this vision, creating the Robot Fabricator system in which physical individuals can be fabricated autonomously. This is a challenging task, but the Robot Fabricator operates within the constraints of current technologies, to create a feasible system for the automatic production of robot bodies, while maximising the diversity of possible morphologies. Automated manufacture of evolved robots in the real world will allow us to address interesting and important questions around morphological evolution in hardware, the reality gap, and how these systems can be implemented, with potential for fundamental advances in the field.

## Acknowledgements

The work reported in this paper is funded by EPSRC under the ARE project: EP/R03561X, EP/R035679, EP/R035733.

## References

- Bartlett, N. W., Tolley, M. T., Overvelde, J. T., Weaver, J. C., Mosadegh, B., Bertoldi, K., Whitesides, G. M., and Wood, R. J. (2015). A 3d-printed, functionally graded soft robot powered by combustion. *Science*, 349(6244):161–165.
- Bongard, J. and Lipson, H. (2004). Once more unto the breach: Co-evolving a robot and its simulator. In *Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems (ALIFE9)*, pages 57–62.

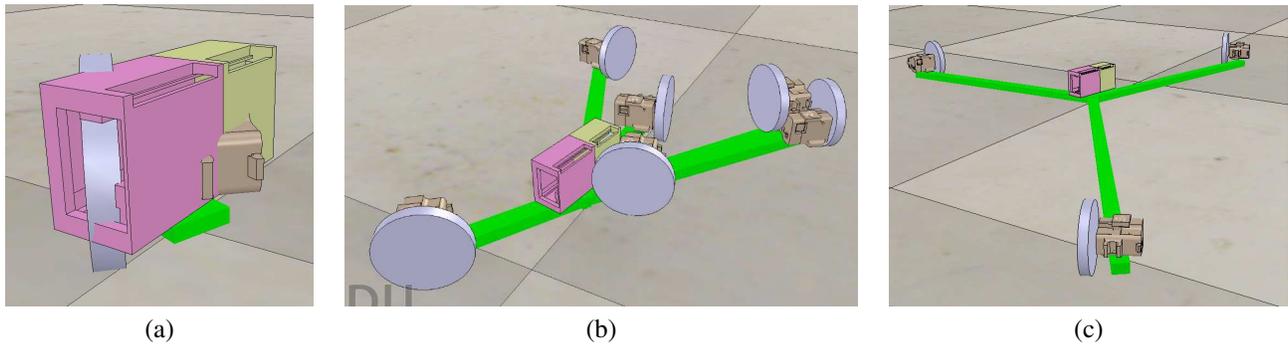


Figure 10: Examples of robots evolved without viability test: (a) robot with overlapping organs, (b) robot with high number of organs and (c) robot bigger than build volume.

- Brodbeck, L., Hauser, S., and Iida, F. (2015). Morphological evolution of physical robots through model-free phenotype development. *PLoS one*, 10(6):e0128444.
- Doncieux, S., Bredeche, N., Mouret, J.-B., and Eiben, A. (2015). Evolutionary robotics: what, why, and where to. *Frontiers in Robotics and AI*, 2:4.
- Eiben, A. (2015). EvoSphere: The World of Robot Evolution. In Dediu, A.-H., Magdalena, L., and Martín-Vide, C., editors, *Proc. of the Theory and Practice of Natural Computing 2015*, LNCS 9477, pages 3–19. Springer.
- Eiben, A., Bredeche, N., Hoogendoorn, M., Stradner, J., Timmis, J., Tyrrell, A., and Winfield, A. (2013). The triangle of life: Evolving robots in real-time and real-space. In Lio, P., Miglino, O., Nicosia, G., Nolfi, S., and Pavone, M., editors, *Proc. of the 12th European Conference on the Synthesis and Simulation of Living Systems (ECAL 2013)*, pages 1056–1063. MIT Press.
- Eiben, A., Kernbach, S., and Haasdijk, E. (2012). Embodied artificial evolution: Artificial evolutionary systems in the 21st century. *Evolutionary Intelligence*, 5(4):261–272.
- Eiben, A. and Smith, J. (2015). From evolutionary computation to the evolution of things. *Nature*, 521(7553):476.
- Hiller, J. and Lipson, H. (2012). Automatic design and manufacture of soft robots. *IEEE Transactions on Robotics*, 28(2):457–466.
- Hornby, G. S., Lohn, J. D., and Linden, D. S. (2011). Computer-automated evolution of an x-band antenna for nasa’s space technology 5 mission. *Evolutionary Computation*, 19(1):1–23. PMID: 20583909.
- Hwangbo, J., Lee, J., Dosovitskiy, A., Bellicoso, D., Tsounis, V., Koltun, V., and Hutter, M. (2019). Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26):eaau5872.
- Jakobi, N., Husbands, P., and Harvey, I. (1995). Noise and the reality gap: The use of simulation in evolutionary robotics. In *European Conference on Artificial Life*, pages 704–720. Springer.
- Jelisavcic, M., De Carlo, M., Hupkes, E., Eustratiadis, P., Orłowski, J., Haasdijk, E., Auerbach, J. E., and Eiben, A. E. (2017). Real-world evolution of robot morphologies: A proof of concept. *Artificial life*, 23(2):206–235.
- Lipson, H. and Pollack, J. B. (2000). Automatic design and manufacture of robotic lifeforms. *Nature*, 406(6799):974.
- MacCurdy, R., Katzschmann, R., Kim, Y., and Rus, D. (2016). Printable hydraulics: A method for fabricating robots by 3d co-printing solids and liquids. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3878–3885. IEEE.
- Nolfi, S., Bongard, J., Husbands, P., and Floreano, D. (2016). Evolutionary robotics. In *Springer Handbook of Robotics*, pages 2035–2068. Springer.
- Nolfi, S., Floreano, D., and Floreano, D. D. (2000). *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. MIT press.
- Pfeifer, R. and Bongard, J. (2007). *How the body shapes the way we think: a new view of intelligence*. MIT press.
- Pollack, J. B. and Lipson, H. (2000). The golem project: Evolving hardware bodies and brains. In *Proceedings. The Second NASA/DoD Workshop on Evolvable Hardware*, pages 37–42. IEEE.
- Sims, K. (1994). Evolution of Virtual Creatures. *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*.
- Winfield, A. and Timmis, J. (2015). Evolvable robot hardware. In Trefzer, M. and Tyrrell, A., editors, *Evolvable Hardware*, pages 331–348. Springer.
- Zykov, V., Mytilinaios, E., Adams, B., and Lipson, H. (2005). Self-reproducing machines. *Nature*, 435(7039):163–164.
- Zykov, V., Mytilinaios, E., Desnoyer, M., and Lipson, H. (2007). Evolved and designed self-reproducing modular robotics. *Robotics, IEEE Transactions on*, 23:308 – 319.