Hardware Design for Autonomous Robot Evolution

Matthew F. Hale*, Mike Angus†, Edgar Buchanan†, Wei Li†, Robert Woolley†, Léni K. Le Goff‡, Matteo De Carlo§, Jon Timmis¶, Alan F. Winfield*, Emma Hart‡, Agoston E. Eiben§ and Andy M. Tyrrell†
*Bristol Robotics Laboratory, University of the West of England, Bristol, UK
†Department of Electronic Engineering, University of York, UK
‡School of Computing, Edinburgh Napier University, UK
§Department of Computer Science, Vrije Universiteit Amsterdam, NL
¶University of Sunderland, UK
email: matt.hale@brl.ac.uk

***Abstract*—The long term goal of the Autonomous Robot Evolution (ARE) project is to create populations of physical robots, in which both the controllers and body plans are evolved. The transition for evolutionary designs from purely simulation environments into the real world creates the possibility for new types of system able to adapt to unknown and changing environments. In this paper, a system for creating robots is introduced in order to allow for their body plans to be designed algorithmically and physically instantiated using the previously introduced Robot Fabricator. This system consists of two types of components. Firstly, *skeleton* parts are created bespoke for each design by 3D printing, allowing the overall shape of the robot to include almost infinite variety. To allow for the shortcomings of 3D printing, the second type of component are *organs* which contain components such as motors and sensors, and can be attached to the skeleton to provide particular functions. Specific organ designs are presented, with discussion of the design challenges for evolutionary robotics in hardware. The Robot Fabricator is extended to allow for robots with joints, and some example body plans shown to demonstrate the diversity possible using this system of robot generation.**

***Index Terms*—evolutionary robotics, evolution in hardware**

I. Introduction

This paper forms part of the Autonomous Robot Evolution (ARE) project[1], the goal of which is the creation of a robot ecosystem that reproduces and evolves in real-time and real-space, meaning that individual robots are physically created and tested, with the results feeding into the evolutionary algorithm to dictate the physical nature of future generations. Such a system offers the prospect of populations of robots which can be deployed into previously unknown or changing environments, and adapt both in control and physical attributes to achieve a given task.

In biology, the simultaneous evolution of both bodies and brains has lead to highly complex organisms and embodied intelligence which is specific to particular *body plans* (the physical layout of the body parts), with tight interlinking between body and brain [1]. It is also true that for (evolved) robots, the interaction of both controller and body plan (also known as the morphology) with the external environment will determine its behaviour and, ultimately, ability to perform a task or tasks (*fitness*). Therefore, for artificial evolution to approach the complexity and diversity of solutions from natural evolution, the evolution of the body plan as well as the controller for the robots is likely to be the way forward.

Much of evolutionary robotics in hardware has thus far evolved only the controller, in a fixed physical robot without any morphological differences between individuals [2]. This makes the process significantly easier, because each new individual needs only new software to be downloaded, without changing the physical robot. The question of how to make evolvable hardware for robots has been explored much less, but offers much greater scope for approaching the complexity of natural evolution, and is the main topic of this paper.

Evolving the bodies of physical robots poses significant practical challenges. In contrast to evolving only the controller, where the creation of a new individual requires new code to be uploaded to a fixed physical body, for morphological evolution a new robot body must be created. 3D printing offers exciting potential to create many unique individuals from the designs created by evolution, at a reasonable cost and speed. However, current 3D printing technology would not be able to produce components such as control electronics, wires, motors and sensors. Adding these components by hand after printing is a possibility [3], however will become impractical for the numbers of individuals needed for embodied evolution. To make the production of large numbers of individuals simpler, it is common to use a modular system in evolutionary robotics, with prefabricated components which connect directly to one another. However, the limited ways in which modules can be combined leads to a highly restricted search space.

Instead, in this paper it is proposed that a new approach is taken, with the components split into two categories. Firstly, those made specifically for a particular robot to give the overall shape which we call the *skeleton*. These are rigid parts with no actuation or sensing, and so can be produced by readily available and affordable 3D printers, which also allows for automation of their fabrication at relatively low cost. Secondly, components which can be re-used for many individuals, even if these individuals have different body plans. This is akin to modular components in biology, such as a finger, or an eye, which are very similar between related species. As such, these components will be more generalised in comparison to the

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

[1]See www.york.ac.uk/robot-lab/are/

individual specific skeleton parts. The production cost for these components is shared across many robots, so these parts can be hand designed and include the electromechanical components needed for actuation and sensing. Because these components are somewhat self-contained systems which provide a particular function to the robot, the biological analogy is used in this work and they are referred to as *organs*. This division between skeleton and organ creates a framework which allows for an almost infinite array of possible shapes for the body by 3D printing the skeleton, with those functional actuators and sensors which cannot practically be printed contained within organs, which can be attached anywhere onto the skeleton to form the final robot.

The contributions of this paper are both the overall concept for this way of expressing robots so that complex and varied body plans can be evolved, and secondly the particular designs created using this concept for the ARE project. The assembly process for these robots is automated through the *Robot Fabricator*, introduced previously [4], and in this paper the assembly procedure is extended to allow for joint organs and multiple skeleton parts. Combining the robot components described in this paper with the Robot Fabricator creates a system to allow, for the first time, physical robot phenotypes to be used in the evolution of not only the controllers, but also the body plans of robots. As such, this paper seeks to begin to answer the question of how to evolve physical robots and provide a path towards "fully embodied evolution in real-time and real-space" envisioned as the long-term goal [5].

II. Related Work

Evolutionary algorithms have been applied many times to the design of simulated robots (*e.g.* [6]). The desire to physically instantiate evolved robots is motivated by the problem of the reality gap [7]; because any simulation inevitably falls short the exact recreation of the real world, and evolution has a tendency to exploit any difference, a robot designed only in simulation may not transfer well to reality. In the last few years, there has been an exciting move towards systems which aim to create evolved robotic body plans using physical hardware, rather than limiting evolution to only modifying a robot's controller, or only simulating morphological changes.

The first breakthrough in this direction was the Golem project [8], in which an evolutionary algorithm could determine the length of links, which then were custom made for the design using 3D printing and assembled by hand. As such, unique body plans for each individual could be made in the physical world, although the evolution only took place in advance in the simulator, without physical performance being used to compute fitness. The resulting robots were also very simple, with a single type of actuator, no sensors and requiring external power, the only task they could achieve was simple locomotion.

In general, creating a larger morphological search space (*i.e.* a more interesting and challenging evolutionary problem) comes at the expense of complexity in the mechanical design. In particular, modules that can be disassembled and reconfigured tend to create a much more restricted problem, compared with a system able to use bespoke parts for each robot [9].

Because of the practical benefits, several modular systems for evolution in hardware have been created, whereby all the parts of a robot are designed to be easily reconfigured into another design. For example, Faina *et. al.* proposed a set of heterogeneous modules explicitly designed for evolvability [10]. Another example is the RoboGen system [11], which allows evolution of robot body plans using an in-built simulator, but the modules are also designed to be produced in hardware to create physical robots. This system of 3D-printed and hand-assembled modules has been used to create a proof of concept of a single reproductive cycle of physical robots [12]. This consisted of an initial population of two hand-designed robots being evaluated in hardware, and then creating an offspring robot based on a crossover of these two, which was also physically created and placed into the arena. Although highly simplified, this is an important step towards evolving robot populations.

To avoid the reality gap, evolution must take place using evaluations of physical robots. This can be made much more feasible if the individuals to be tested can be constructed automatically. An example of this is the gluing together of simple modules by the "mother robot" [13]. This work is an impressive example of evolution based on fitness values obtained from physical testing, but was limited in scope by very simple robots, with no sensors and a single type of simple actuator.

Another modular system aiming for fully automated assembly of new individuals by an external robot is the "EMERGE" project [14]. Here cube shaped modules snap together magnetically to allow for simple assembly and disassembly by an external robotic arm.

Another approach, to allow for more capable robots, it to have some aspects of the robot hardware which can change or re-configure, such as modifiable leg lengths [15]. This has the advantage of reducing the time and cost to create a new "individual" robot for the population, because it does not need to be created from scratch but rather the existing robot can modify its geometry to the new design. However, the amount of variation possible is very limited, and in particular is limited to those aspects chosen in advance to be adjustable. This is likely to mean that a key advantage of designing robots by evolution is lost: that unforeseen and surprising solutions may be found by evolution that a human may never have considered.

III. Design Objectives and Challenges

Designing a system of components to allow evolution of robot body plans creates some particular requirements and challenges for the design. Furthermore, creating large numbers of individual robots will be facilitated by an automated assembly process (in the Robot Fabricator [4]), which places further restrictions on the design. This section will outline the main challenges that have so far been encountered in the ARE

project, and the next section details the designs created in light of these.

Firstly, the system for the robots must allow for a wide range of varied and complex body plans. Otherwise, the search space for the evolutionary algorithm will be small which will lead to uninteresting results and fail to demonstrate the exploratory power of the algorithms. Instead, a system which allows a wide diversity of body plans could lead to interesting, even unexpected, solutions to a particular task.

On the other hand, if each body plan generated were completely unique, the time and resources to produce each individual would become prohibitive, given the numbers required for evolution. In order to share the cost (in resources and time) of electromechanical components such as motors, sensors, batteries etc., these should be built into organs so they can be reused between individuals. In this way, one generation of robots can be deconstructed in order to assemble the next. This means these components much be attached together via a non-permanent mechanism, that nonetheless provides the strength and rigidity needed for them to perform their function.

To allow for complete populations to be produced in a lab setting, the time taken for production of any custom parts and the assembly of an individual should be minimised. A long production time for each individual would severely reduce the practicality of evolution in hardware, given the numbers required for a complete evolutionary run, even if a hybrid physical-simulation system is implemented.

Large numbers of robot individuals also requires that the components can each me made low-cost. Even if organs can be reused between individuals, it will be important to be able to have several robots constructed simultaneously to allow parallel testing, and the number of each type of component needed for each individual cannot be known in advance (as this is selected by the evolutionary algorithm), and so even a modest population of physical robots will need at least dozens of the various organ types to be made.

Many of the design challenges are imposed by the desire to allow for autonomous assembly in the Robot Fabricator. The mechanical connection between the different components is one such key area. This must be strong and stiff enough for the robot (of unknown configuration) will function without breaking apart, but autonomous assembly precludes fixings such as screws.

As well as mechanical connections, the organs must be electrically connected to for the central controller to receive sensor values and dictate actuator movement. Also sharing power between organs will allow for a simplified power system overall with a centralised battery, also simplifying the logistics of recharging the robot and managing power across multiple organs.

All the organs must also be restricted in size and shape so they can be picked up by the gripper of the Robot Fabricator, which must also be able to handle the cable connections. Using a custom gripper has allowed it to be designed around the task, but adds more complexity to the design.

Upon applying a novelty search algorithm in simulation, another impact of the design of the organs became apparent. Although the novelty search should use each organ type equally, early results showed fewer wheel organs compared to the others. Further experiments showed this was because of the original shape of this organ, with the clip at $90°$ to the wheel, which frequently creates a collision with the skeleton to which it is attached, resulting in the organ being discarded. Instead, placing the clip in line with the wheel prevents this happening and increases the resulting frequency of wheel organs. This demonstrates the potential for seemingly unrelated design decisions to impact upon the final evolutionary process, which may then become biased toward certain solutions by the components available.

IV. Organ designs

A. Cable connections between organs

Each of the peripheral organs need to be connected to the controller located in the head organ for control purposes, providing sensor readings or receiving values for actuator outputs. This communication could be achieved through a wireless system, with all the organs within a robot connected together using WiFi, ZigBee or similar. However, this adds complexity to individual organs, each requiring an independent battery (with associated charging) and wireless hardware, and to the overall system, requiring network setup and management. Overall, this will create many potential points of failure, with a single organ dropping from the wireless network or running out of charge likely ruining an experiment. For this reason, the choice here, at least initially, is to connect all the organs with wires. This also allows for the sharing of power, with all organs powered by a central battery in the head. Power (ground and $5\,\mathrm{V}$) and communications (via an i2c bus) is carried in four wires, with connections made using $3.5\,\mathrm{mm}$ TRRS sockets. These sockets have been chosen because they are readily available, being commonly used for audio with a microphone channel, and can be inserted in any angle around their axis, so a $90°$ jack can be pointed in the desired direction when inserted.

Initially, straight cables were used for the connections. However it soon became clear these would create tangles, with the first cables inserted becoming an obstruction to the assembly of later ones. This problem can be mitigated using coiled cables with some tension applied between the organs to be connected, as shown in Fig. 1. In this way, the cable will form approximately a straight line, being more compact and less likely to cause problems by becoming tangled. They also become more predictable, and if it is found to be necessary in the future it could be possible to automatically predict and avoid cable configurations which would create tangles.

B. Head organ

The core of any robot is a power supply and electronic control. These functions are required for each individual robot produced, so for simplicity we have limited the design to a single, centralised battery and voltage regulation and

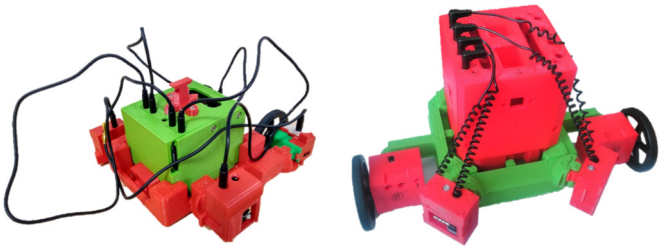

(a) (b)

Fig. 1: Early physical iterations of the organs used to create some hand-designed test robots, showing (a) the tangle-prone cables, and (b) the self-retracting, coiled cables used to reduce this problem.

Raspberry Pi based "brain". These parts will be contained in a single organ called the head, shown in Figs. 2 and 3, and since each individual will include exactly one of these it is used to define a coordinate system for the robot, *i.e.* a centre point and forward direction (although the robot is free to move in any direction it is able to).

The head includes eight female sockets for cables to be inserted into in order to form connections to the peripheral organs. The extendable cable from each will be inserted into one of these sockets, creating as direct as possible a route to the power supply, as discussed above and shown in Fig. 1b.

During assembly, it is also useful to have a defined central point in the robot, and the head must be held securely by the assembly fixture while the rest of the robot is assembled around it. This is accomplished by the inclusion of a set of ferrous metal plates on the bottom of the head organ which interface with electromagnets on the assembly fixture. This system allows a firm mechanical connection during assembly but easy detachment by deactivating the electromagnets once the robot assembly is complete.

The attachment between the head organ and the skeleton is a critical part of the design, and is shown in Fig. 4. In the planned autonomous assembly process, the head is inserted into the appropriate skeleton part when the skeleton has just been created on the 3D printer, and is then used as a handle to allow the robot arm to remove skeleton from the printer bed. When this part of the process was done manually [4], we found it often takes considerable force, and so the head to skeleton connection needs good strength. In the next part of the assembly process, the head is held in the assembly fixture as other organs are attached to the skeleton; in order for the connection points for these peripheral organs to be predictably located, there must be minimal movement between the head and skeleton, so this connection must also be stiff. Furthermore, the design is complicated by the fact that the skeleton must be 3D printed, and to minimise the time taken to create each individual we have selected a printing setup[2]

[2]The printer being used is the Lulzbot TAZ 6, with a 1.2 mm nozzle and 0.9 mm layer thickness.

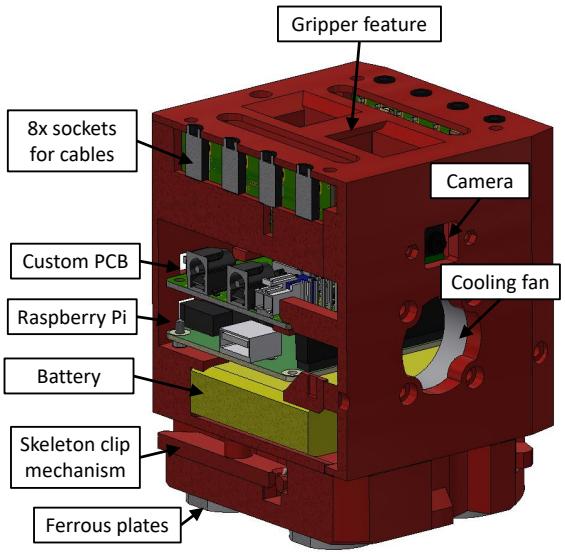

Fig. 2: Section through the head organ design, showing the key components. The head forms the core of the robot, providing power and control for the other organs, which are connected via cables to the sockets on the top.

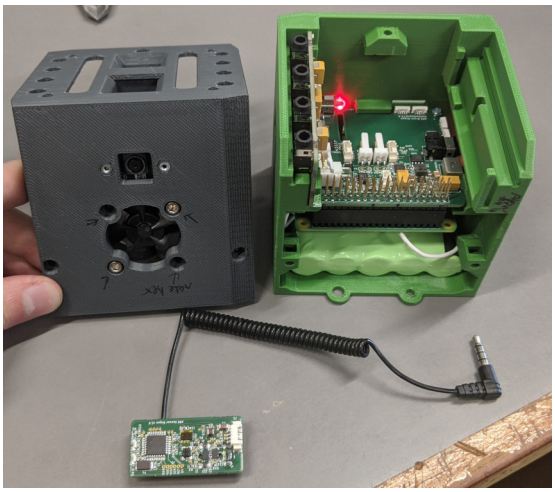

Fig. 3: Head organ prototype, undergoing testing with the PCB from a sensor organ.

which prevents fine detail or high resolution. Any parts that are not 3D printed or require intricate geometry must only be used on the organ side of the clip, not the skeleton side. The resulting design is shown in Fig. 4, with four protrusions of the skeleton preventing the head from lifting out of the skeleton. The head is inserted by a rotating motion, with two sprung levers latching into place against indents in the skeleton.

C. Sensor organ

To perform useful or complex tasks and/or operate in an unknown environment, robots must sense the world around them. Many practical tasks such as exploring, collecting samples or emergency response can be abstracted into some

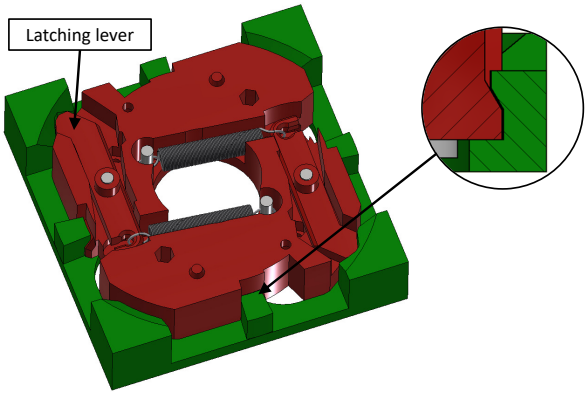

Fig. 4: The clip mechanism which attaches the head organ (red parts) to the skeleton (green). The head is inserted at an angle and then rotated, causing the latching levers to lock it into place.

combination of maze solving (moving around an environment while avoiding obstacles) and foraging or beacon following (moving towards or interacting with a target). So, two separate sensors are provided for these two types of functions. To avoid walls or obstacles, a laser rangefinder[3] has been selected which provides information on the distance to the nearest object in a particular direction. To create a target object which can be differentiated from obstacles, an infrared (IR) beacon will be used, with a simple IR sensor[4] able to detect its presence in a particular direction. Detecting the beacon can be done with the sensor in a passive mode, without the emitter powered.

For practical reasons for the Robot Fabricator, it is desirable to minimise the number of unique physical organs. The decision was made to combine the different sensors into a single physical organ with the sensors side-by side, the design of which is shown in Figs. 5a and 6a. It will be straightforward, in software, to restrict the controller to only using one if a particular sensor has been selected by evolution.

D. Wheel organ

A wheel organ has been chosen as the first actuation organ, because they are expected to be the most straightforward way for a robot to move around its environment. They should provide good efficiency and simple control so long as the floor is smooth and flat enough for the wheels to function. The organ, shown in Figs. 5b and 6b, consists of a pre-made wheel, a 3D printed case incorporating the female half of a clip for attachment to the skeleton and a expendable cable as described above. Internally, a custom PCB houses a motor driver and microprocessor, which reads the wheel encoder and applies a PID controller to achieve the demanded wheel velocity.

E. Joint organ

The design of the joint organ, shown in Figs. 5c and 6c, is based around a servo motor[5] which gives closed loop position control and good torque at low cost in a convenient package. It provides a more complex form of actuation compared to the wheel, and should allow for more complex robot body plans, as discussed below. By including both a wheel and joint organ in the system, the hope is to be able to make comparisons between the evolution of wheeled and limbed robots, and perhaps discover some interesting designs which combine the two.

A joint, by definition, creates movement of two parts of the body relative to each other, so this organ includes two clips, one each side of the hinge.

V. Limb Assembly

To allow for autonomous evolution, the robots must have their skeleton created and then be assembled from the various components autonomously. This has been the source of many of the design constraints described in previous sections, and is closely linked to the clip design for connecting the organs to the skeleton. The concept of the Robot Fabricator, a machine to achieve this, was introduced in a previous paper [4]. However, in that paper we considered the assembly of simple robots, where a single skeleton part, around the head, has all the other organs attached directly to it. This was achieved by mounting the head and skeleton into a central assembly fixture, clipping each organ on in turn.

Here the process is extended by the introduction of the joint organ, which creates the necessity for multiple skeleton parts attached to each side of the articulating joints. Fig. 7 shows conceptually a jointed robot body plan, formed of alternating skeleton parts (green) and organs (red) with the possibility to form a branching tree structure with multiple organs connected to a skeleton part. Every robot must have a head organ, and some skeleton attached directly to it ("skeleton 0" in the figure). Each joint that is then added to the skeleton allows a further skeleton part to be added to the other side of the joint, so that the total number of skeleton parts is always one more than the number of joints in the robot. Other organs, such as wheels and sensors, can be added to any piece of skeleton. The alternation between skeleton and organs means the skeleton needs only to include the male half of the clip, which has been designed to be easily 3D printed. A Joint organ and everything connected to in (possibly including another joint) can be considered as a "limb" of the robot.

Assembly of these limbs poses a challenge to the automated assembly process proposed previously [4]. In this process, the head is held in an assembly fixture, with other parts progressively added to it. Consider this situation when "skeleton 2" must be attached to "joint 2". There are already three clips between this one and the head organ, which is the only part firmly fixed. Inevitably, each clip will introduce some inaccuracy and wobble, and connecting them in series causes these inaccuracies to compound. This is a problem known as tolerance stack up, and will mean that the attachments at the end of a long limb will become too difficult for the fabricator.

To overcome this problem, a proposed scheme is shown in Fig. 8, whereby the limb is gradually built up from its

[3]STMicroelectronics VL53L0X Time-of-Flight Ranging Sensor
[4]Sharp GP2S700HCP infrared sensor
[5]The Towerpro MG996R

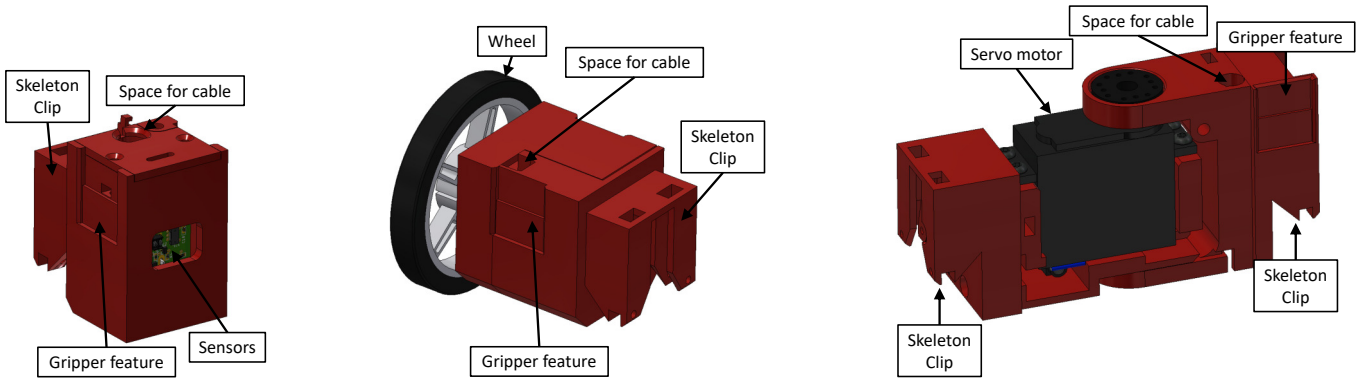

(a) (b) (c)

Fig. 5: Designs for the (a) sensor, (b) wheel and (c) joint organs.

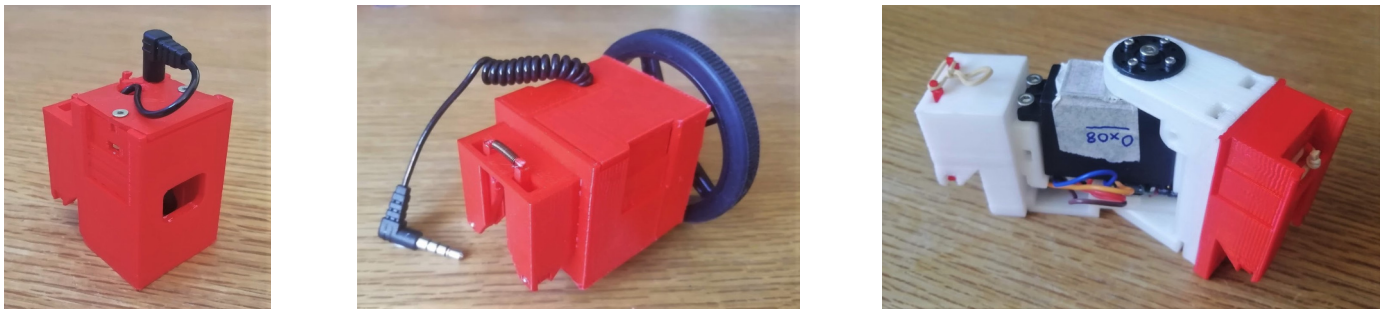

(a) (b) (c)

Fig. 6: Physical prototypes of the (a) sensor, (b) wheel and (c) joint organs.

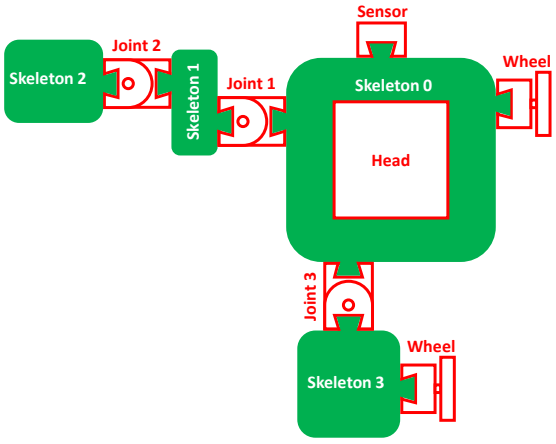

Fig. 7: A conceptual layout for for skeleton and organs might be combined to form a robot. The "skeleton 0" is connected directly to the head, with other skeleton parts attached the other side of joints. Because the skeleton includes the male side of the clip and the organs the female, they must be connected in a tree structure, alternating between skeleton and organ.

extremity first, so that the entire limb can then be attached as a single part to "skeleton 0". The process is designed so that when each part is connected to its parent in the robot tree structure, its parent is held directly in a fixture in order to minimise the tolerance stack-up. The process begins (1) with a fixture containing dummy versions of the male and one female sides of the clip, which will hold the parts in place but allow the robot arm to remove them later. The assembly begins with the parts which will be furthest from the head, attaching a joint to the fixture and the last skeleton part to it (2-3). Then the appropriate skeleton part is attached to the female half of the fixture (4) and the joint is moved across onto it (5). The process of adding a new component to the fixture, then moving the partly assembled limb onto it continues (6) until the limb is complete (7), at which point it can be attached to the robot on the main assembly fixture.

This process has been demonstrated on the existing Robot Fabricator using a hand-designed example limb, as shown in Fig. 9. The process is as described above, but only a single joint has so far been implemented for this demonstration.

VI. Example Body Plans

By way of demonstration that the organ-skeleton paradigm presented here is able to produce a diverse and interesting range of body plans, this section will briefly present some such body plans made made by two methods.

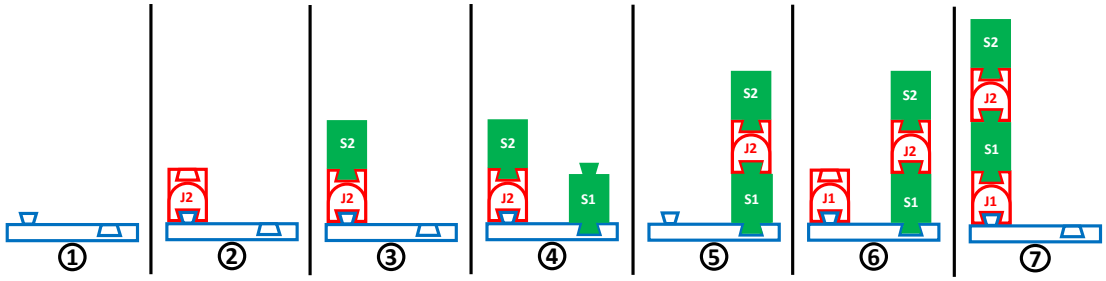

Fig. 8: Conceptual diagram of how a complex limb can be build up on a fixture, where each attachment occurs onto a component which is rigidly held. "J0/1" corresponds to "Joint 0/1" and "S0/1" to "Skeleton 0/1" from Fig. 7.

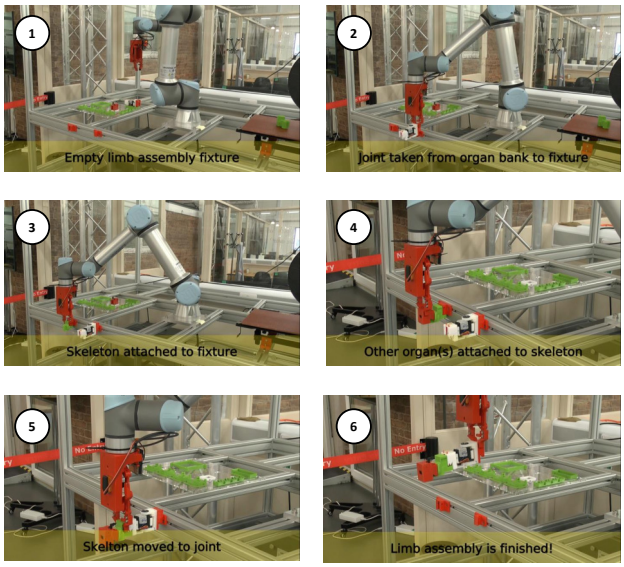

Fig. 9: Key steps of the demonstration limb assembly. The full video is available at https://youtu.be/6D6nMtGryME

In advance of the morphological evolution being implemented, some hand-designed body plans have been created, shown in Fig. 10. These body plans are being used to test and develop some of the hardware, such as the Robot Fabricator [4], and software, such as selecting and developing the controller architecture and learning algorithms which will eventually be applied to evolved robots [16]. They are therefore designed to require different controllers, with a maze solving task in mind. The *ARE-puck* is designed to be simple to control for tasks such as maze solving, with plenty of sensors and two wheels arranged symmetrically on each side. Adding some more difficultly based on expectations of what evolution may produce, the *Potato* has reduced sensing capability and asymmetrical wheels. Finally, the *Tricycle* is a very different type of robot, utilising both a joint and a wheel. It is designed so that the wheel can drive the robot forward, with steering achieved by the joint pivoting the wheel.

However, manually designed robots are not the end goal of this project, but rather a diverse range of automatically generated morphologies, such as those shown in Fig. 11. These morphologies have been generated by novelty search, and show a range of morphologies can be found by an evolutionary algorithm using these organs [17].

VII. Conclusions

The creation of physical robots which allow for the evolution of body plans requires a system of construction which can produce a wide range of designs able to be encoded into a genome. This motivates a highly modular approach, so that functional parts can be rearranged to create different body plans. However, such modules do not allow for as much flexibility in design as custom made parts; producing bespoke parts for each individual can be enabled by 3D printing to give an almost infinite range of possible geometries, yielding a large and interesting search space for evolution to explore.

The vision of the Autonomous Robot Evolution project is to create a population of evolving robots in hardware. This paper is a step towards that goal, combining reusable organs with an individual specific skeleton made by 3D printing. There are many challenges to designing the system and in particular the organs, many of which become apparent only upon implementing the system, especially in order to allow for the automatic assembly of the resulting robots.

The example body plans presented here have yet to be evolved for a task, but they do show the diversity of potentially functional robots this system can generate. Closing the evolutionary loop in simulation is a priority for future work, while for the physical hardware the next steps will be to create multiple physical organs to allow for robots to be created and tested in the real world.

References

[1] R. Pfeifer and J. Bongard, *How the body shapes the way we think: a new view of intelligence*. MIT press, 2007.

[2] S. Nolfi, D. Floreano, and D. D. Floreano, *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. MIT press, 2000.

[3] J. B. Pollack and H. Lipson, "The golem project: Evolving hardware bodies and brains," in *Proceedings. The Second NASA/DoD Workshop on Evolvable Hardware*. IEEE, 2000, pp. 37–42.

[4] M. F. Hale, E. Buchanan, A. F. Winfield, J. Timmis, E. Hart, A. E. Eiben, M. Angus, F. Veenstra, W. Li, R. Woolley *et al.*, "The are robot fabricator: How to (re) produce robots that can evolve in the real world," in *Artificial Life Conference Proceedings*. MIT Press, 2019, pp. 95–102.

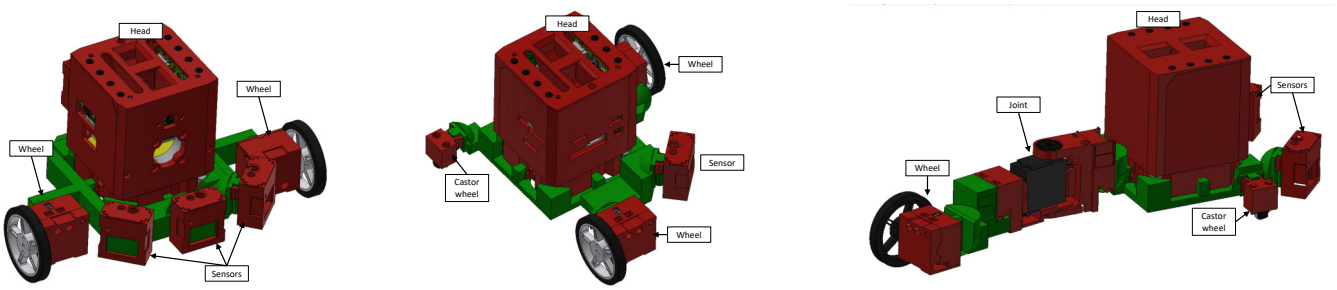

(a) (b) (c)

Fig. 10: Example body plans, designed by hand to demonstrate a range of possible robots for use in testing the physical assembly process and the simulation environment. (a) is the *ARE-puck*, which is symmetrical and designed to be easy to control, (b) the *Potato* which is similar but asymmetrical to add some complexity, and (c) the *Tricycle* which requires control of a joint and a wheel.

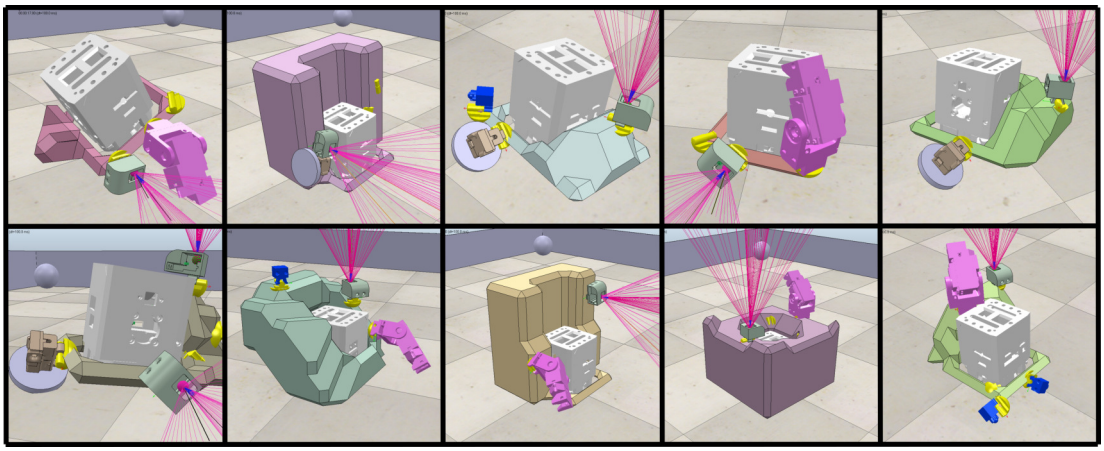

Fig. 11: A range of possible body plans created by novelty search, without any consideration for task performance or locomotion.

[5] A. Eiben, N. Bredeche, M. Hoogendoorn, J. Stradner, J. Timmis, A. Tyrrell, and A. Winfield, "The triangle of life: Evolving robots in real-time and real-space," in *Proc. of the 12th European Conference on the Synthesis and Simulation of Living Systems (ECAL 2013)*, P. Lio, O. Miglino, G. Nicosia, S. Nolfi, and M. Pavone, Eds. MIT Press, 2013, pp. 1056–1063.

[6] K. Sims, "Evolution of Virtual Creatures," *Proceedings of the 21st annual conference on Computer graphics and interactive techniques.*, 1994.

[7] N. Jakobi, P. Husbands, and I. Harvey, "Noise and the reality gap: The use of simulation in evolutionary robotics," in *European Conference on Artificial Life*. Springer, 1995, pp. 704–720.

[8] H. Lipson and J. B. Pollack, "Automatic design and manufacture of robotic lifeforms," *Nature*, vol. 406, no. 6799, p. 974, 2000.

[9] R. Moreno and A. Faina, "Reusability vs morphological space in physical robot evolution," in *Proceedings of the 2020 Genetic and Evolutionary Computation Conference Companion*, 2020, pp. 1389–1391.

[10] A. Faina, F. Bellas, F. Orjales, D. Souto, and R. J. Duro, "An evolution friendly modular architecture to produce feasible robots," *Robotics and Autonomous Systems*, vol. 63, pp. 195–205, 2015.

[11] J. Auerbach, D. Aydin, A. Maesani, P. Kornatowski, T. Cieslewski, G. Heitz, P. Fernando, I. Loshchilov, L. Daler, and D. Floreano, "Robogen: Robot generation through artificial evolution," in *Artificial Life Conference Proceedings 14*. MIT Press, 2014, pp. 136–137.

[12] M. Jelisavcic, M. De Carlo, E. Hupkes, P. Eustratiadis, J. Orlowski, E. Haasdijk, J. E. Auerbach, and A. E. Eiben, "Real-world evolution of robot morphologies: A proof of concept," *Artificial life*, vol. 23, no. 2, pp. 206–235, 2017.

[13] L. Brodbeck, S. Hauser, and F. Iida, "Morphological evolution of physical robots through model-free phenotype development," *PloS one*, vol. 10, no. 6, p. e0128444, 2015.

[14] R. Moreno, F. Veenstra, D. Silvera, J. Franco, O. Gracia, E. Cordoba, J. Gomez, and A. Faina, "Automated reconfiguration of modular robots using robot manipulators," in *2018 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE, 2018, pp. 884–891.

[15] T. F. Nygaard, D. Howard, and K. Glette, "Real world morphological evolution is feasible," *arXiv preprint arXiv:2005.09288*, 2020.

[16] L. K. Le Goff, E. Buchanan, E. Hart, A. E. Eiben, W. Li, M. de Carlo, M. F. Hale, M. Angus, R. Woolley, J. Timmis *et al.*, "Sample and time efficient policy learning with cma-es and bayesian optimisation," in *Artificial Life Conference Proceedings*. MIT Press, 2020, pp. 432–440.

[17] E. Buchanan, L. K. Le Goff, E. Hart, A. E. Eiben, M. de Carlo, W. Li, M. F. Hale, M. Angus, R. Woolley, A. F. Winfield, J. Timmis, and A. M. Tyrrell, "Evolution of diverse, manufacturable robot body plans," in preparation.