Proceedings Paper:
Buchanan Berumen, Edgar, Le Goff, Léni, Hart, Emma et al. (Accepted: 2020) Evolution of Diverse, Manufacturable Robot Body Plans. In: 2020 IEEE Symposium Series on Computational Intelligence. (In Press)

Evolution of Diverse, Manufacturable Robot Body Plans

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

***Abstract*—Advances in rapid prototyping have opened up new avenues of research within Evolutionary Robotics in which not only controllers but also the body plans (morphologies) of robots can evolve in real-time and real-space. However, this also introduces new challenges, in that robot models that can be instantiated from an encoding in simulation might not be manufacturable in practice (due to constraints associated with the 3D printing and/or automated assembly processes). We introduce a representation for evolving (wheeled) robots with a printed plastic skeleton, and evaluate three variants of a novelty-search algorithm in terms of their ability to produce populations of manufacturable but diverse robots. While the set of manufacturable robots discovered represent only a small fraction of the overall search space of all robots, all methods are shown to be capable of generating a diverse population of manufacturable robots that we conjecture is large enough to seed an evolving robotic ecosystem.**

***Index Terms*—evolutionary robotics, autonomous robot fabrication, autonomous robot evolution, robot manufacturability**

I. Introduction

The Evolution of Things was introduced to describe a new type of Evolutionary Computation that represents a departure from the evolution of digital artefacts to the evolution of physical ones [1], [2]. Advances in robotics, 3D-printing, and automated assembly techniques have recently provided us with the tools required to realise such systems, opening up new avenues of research in which evolution can take place completely in hardware. This is of particular interest to the evolutionary robotics community, as crossing the infamous reality-gap [3] hinders the transfer of robots evolved in simulation into the real-world.

However, while engineering advances in materials, printing and automated assembly offer an unprecedented opportunity to study embodied evolution, and the evolution of ecosystems of physical robots in which bodies and controllers co-evolve, it also introduces significant new challenges that do not appear when only evolving in the virtual world [4]. In order to translate a genotype into a physical phenotype comprising a robot body, sensors, actuators and brain (i.e. software controller), two factors must be considered. Firstly, some components need to be *manufactured* from scratch according to the phenotype (for example, 3D-printing of an evolved skeleton), then secondly, the components need to be *assembled* into the desired phenotype.

A novel instantiation of a system to realise this was proposed by Hale *et al* [5], depicted in Figure 1, developed as part of the ARE project[1] which consists of a 3D printing station used to create a skeleton, combined with a robotic arm which then attaches pre-built components (organs) and finally inserts the necessary wiring. This system illuminates several issues related to manufacturing. Some issues are related to 3D printing, for instance, overhanging sections of the skeleton cannot be 3D printed without the aid of supporting material which is difficult to remove. Additionally, the assembly process introduces its own constraints; for instance, an assembly arm might not be able to manoeuvre into the required position to insert an actuator or sensor. Such issues have not been considered before because they do not appear in simulation; they are rooted in the physical nature of the objects that evolve. This illustrates that the Evolution of Things in general and the evolution of real robots in particular is intrinsically different from evolutionary computation.

The main objective of this paper is to gain insights into how the introduction of constraints associated with manufacture and assembly influences the evolutionary process. In addition, we need to identify the proportion of robots that are manufacturable. With this information we can design better EAs that focus on the manufacturable parts. Specifically, we are concerned with:

- the *manufacturability* of robots, i.e. the ability of a phenotype to be both printed and assembled in an automated process,
- the *viability*[2] of evolved robots, i.e whether they meet a minimal requirement in terms of their actuators/sensors to function usefully, and

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

[1] Autonomous Robot Evolution: https://www.york.ac.uk/robot-lab/are/

[2] In [5], the term viability was used to cover both manufacturability *and* functional considerations: here we separate the two concepts

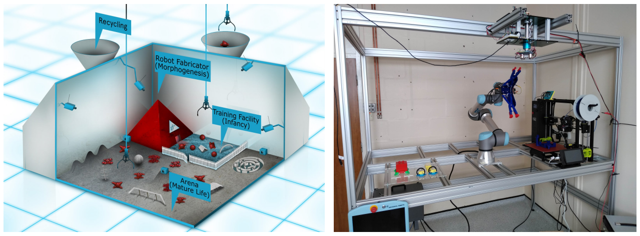

Fig. 1. Illustration of the ARE environment (left), showing the three main stages of the project: Robot Fabricator, Training Facility and Arena. The evolved body plans are manufactured in the Robot Fabricator (right).

- the *diversity* of the evolved population.

The latter may not seem directly related to the physical embedding, but it is important for practical reasons. Given that the search-space of body plans (morphologies) and controllers is very large, and printing/assembly is both time-consuming and expensive, we intend to bootstrap our evolutionary process by starting with a diverse population of manufacturable robots. For instance, it was shown by Le Goff *et al* [6] that for only learning it takes at least a couple hundred of evaluations to produce controllers for ARE-robots to solve tasks.

The system we use and study in this paper is based on an indirect encoding that produces robots comprising plastic skeletons equipped with wheels and sensors, and is used in conjunction with a novelty search [7] algorithm to search for diverse body plans. Importantly, we augment the novelty search algorithm with repair mechanisms to ensure manufacturability of the resulting robots. We distinguish two types of repair mechanism: in-evolution repair and post-evolution repair (explained later on). They are similar in what they do — both overrule the instructions coded in the genotype in order to reduce the number of impossible positions or orientations of 'body parts'— the difference is in *when* they do this. The specific research questions we investigate are the following:

1) How do the methods compare in terms of the proportions of i) manufacturable and ii) viable robots produced?
2) How do the evolutionary methods compare in terms of the diversity of the evolved manufacturable and/or viable population?
3) Is there a trade-off between manufacturability and diversity?

The results demonstrate that the highest diversity of robots is achieved when there are no constraints in the evolutionary process. However, only a small fraction of evolved robots are manufacturable. The repair mechanisms increase the number of manufacturable robots, but they also decrease the diversity. Nevertheless, all algorithms produce populations of at least 175 manufacturable, viable robots, a size that seems appropriate to seed a future evolutionary process.

II. Related Work

Evolving in hardware requires consideration of both *manufacturing* and *assembling* processes. Most existing work within evolutionary robotics focuses on the latter and can be described under the general heading of modular robotics. Here, an EA searches for an arrangement of a fixed set of components (off-the-shelf or printed) that are then assembled following a simulated evolution process. For example, the Golem project [8] evolved robots in a design space comprising of a set of bars and actuators, connected with free joints, which were assembled by hand post-evolution.

The 'Robot Baby' project [9] also followed a modular approach in evolving new body plans from a fixed set of components, although a notable difference to previous work was that reproduction took place in hardware and parents and offspring co-existed in the same arena. However, newly generated robot designs were built by hand.

A move to automated assembly was described by Brodbeck *et al* [10] who again evolved robots using fixed modules but the assembly was completely automated, with a robotic arm picking modules and gluing them together according to evolved designs. The uniformity and simplicity of the components used (passive and active cubic modules) contributed to making the automated assembly relatively straightforward in this case; more complex and diverse components would clearly introduce assembly issues.

Most recently, an EA was used to evolve a robot in simulation which was then built by hand from living cells, creating the world's first living robot (xenobot) [11].

With respect to manufacturing, Funes *et al* [12] evolved modular manufacturable structures capable of support external loads. Collins *et al* [13] use an EA followed by 3D-printing to design a *single* component, specifically a leg for a walking robot, applying repair methods to ensure the resulting designs are manufacturable. Hiller and Lipson [14] introduced a method of evolving robots with only actuators and simple controllers in simulation which were *assembled* automatically. Recent advances in multimaterial fabrication techniques were exploited to evolve freeform soft robots with forward locomotion from soft volumetrically expanding actuator materials; however, as in previous work in modular robotics, the components were uniform in their size and shape, aiding manufacturing.

In our recent work, Hale *et al* [5] outlined the challenges associated with automated manufacture and assembly of evolved designs which include a broad range of components which are diverse in shape and size, as well as diverse skeleton body plans (manufactured from hard-plastic). That paper first introduced the notion of a *manufacturability test*. This article builds directly on this work in understanding how accounting for manufacturability and automated assembly influences the evolutionary process.

III. Experimental Methodology

The genome and the decoder to produce the body plans of the robots produced in this paper are first described in this section. Then, the manufacturability constraints are introduced. Lastly, the algorithms to explore the diversity and manufacturability search spaces are explained.

A. Body plan generation

The robots presented in this paper have two main components: skeleton and organs. The skeleton is the 3D-printable structure to which the organs are attached. The organs are the pre-fabricated (non printable) components of the robots. The organ types used in this study are wheels, sensors and head organ (although joint organs are currently being integrated). A more technical description about the different organ types can be found in [15]. The body plans are encoded in a Compositional Pattern Producing Network (CPPN) firstly introduced by Stanley [16]. The structure of the CPPN consists of four inputs (x, y, z and r) and four outputs. The four inputs define the position of a cell in a 3D matrix, where x, y and z are the coordinates of the cell and r is the distance of the cell from the centre of the matrix. The four outputs define the properties of each cell in the 3D matrix. The first three outputs are binary and they represent the absence or presence of a specific type of voxel. There are three types of voxels: skeleton, wheel and sensor. More types can be added in the future for different organs. The last output represents the rotation of a specific component along the normal of the surface of the skeleton. The genotype to phenotype decoder is as follows:

1) Firstly, the skeleton output is queried for the entire 3D matrix. This will generate the plastic connecting all the organs (passive and active components). Only one region of skeleton is allowed per body plan, where a region is cluster of inter-connecting skeleton voxels. The biggest region of skeleton is preserved and any regions unconnected to the largest region are removed.
2) Then, the CPPN is queried to determine wheel and sensor outputs, generating multiple regions for each type. An organ is generated in each intersecting area between the organ region and the skeleton surface. This ensures that all the components are connected to the surface of the skeleton. Only one component is generated in each region regardless of the size of the intersecting area.
3) Two relative rotations of the organs are given by the normal of the skeleton surface. The third rotation along the relative z-axis of the organ is given by the last output of the CPPN.

B. Manufacturability constraints imposed by the Robot Fabricator

The Robot Fabricator [5] (Figure 1) imposes various constraints on the body plans that can be manufactured. The constraints considered with respect to a robot being declared 'manufacturable' are listed below:

- *Skeleton presence test* - Sometimes a CPPN does not output any skeleton voxels. This is invalid because the organs have nothing to connect to.
- *Skeleton connected to head organ test* - the head organ is fixed in the same position and orientation for all the body plans and the skeleton is generated around it. However, sometimes the skeleton generated is not connected directly to the head organ: this is not valid.
- *Colliding organs test* - Body plans evolved in simulation should not have overlapping components (skeleton or organs).
- *Good organ orientation test* - The organs are attached to the skeleton through male-female clip connectors where the male connector is always generated on the skeleton. The male connectors should not be pointing downwards otherwise supporting material is required from the 3D printer. This supporting material would act as an obstruction.
- *Gripper access test* - The robot arm should have clear access to attach the organs. In other words, there should not be any obstacles in the way of the robot arm.

For the results presented in this paper a robot is considered viable when the body plan has at least 1 wheel and 1 sensor.

With these extra constraints in mind, the next section considers three evolutionary algorithms to create manufacturable and viable physical robots

C. Evolutionary Algorithms

Four different algorithms were used for the experiments described in this paper: Random Sampling (*RS*), Novelty Search (*NS*), internal Gene Repression (*GRi*) and external Gene Repression (*GRe*).

1) Random Sampling (RS):: For the baseline experiments 20,000 random robots are generated for each repetition. This is achieved by randomising the weights and activation functions of the feedforward CPPN described in section III-A.

2) Novelty Search(NS):: For these experiments the aim is to generate a diverse population of robots. We use the novelty search algorithm [7] which replaces the traditional objective function of an EA that rewards objective fitness of an individual with one which rewards novelty w.r.t a user-defined descriptor. It has been shown to both overcome deceptiveness and to be able to locate high-quality solutions in unexpected parts of the search-space. The algorithm maintains an archive of novel solutions, and measures novelty of each solution w.r.t the other members of the current population *and* those maintained in the archive. The reader is referred to [7] for a full description of the algorithm; here we just describe in detail those aspects that are specific to our implementation.

As previously explained, body plans are generated by a CPPN. We use HyperNEAT [17] with its default parameters to evolve a population of CPPNs from which body plans are generated. We measure novelty w.r.t a 6-dimensional morphological descriptor obtained from the phenotype. This contains:

- The *width*, *depth* and *height* attributes describe the volume of the robot.
- The *number of wheels* and *number of sensors* attributes represent the final number of these components in the body plan.
- The *number of voxels* represents the voxels used for the static skeleton in the body plan.

Novelty is calculated according to the sparseness metric defined in equation 1, where k represents the number of closest neighbours in the archive, μ_i represents neighbour i, and the $dist$ function represents the Euclidean distance between two descriptors. A value of 15 for k has been previously shown to provide good results [18] and is used in all experiments.

$$\rho(x) = \frac{1}{k}\sum_{i=0}^{k} dist(x, \mu_i), \tag{1}$$

The NS method requires a method of adding novel solutions to the archive to be defined. Here, descriptors are added to the archive if they meet one of the two following conditions. Firstly, each descriptor x has a probability $P(x) = 0.1$ of being added to the archive automatically. Second, a descriptor is added if the sparseness value $\rho(x)$ is greater than 0.2 ($\rho(x) > 0.2$). Note that manufacturability is not evaluated as part of the NS algorithm.

3) Repair Method in the Evolutionary Loop (GRi):: This algorithm is a variation of NS. Instead of evaluating the robot directly, a correction mechanism (if required) attempts to make the robot manufacturable before it is evaluated.

The correction mechanism used in this paper is referred as *gene repression*. If a single organ fails a manufacturability test, this organ is removed from the final body plan. The morphological descriptor is derived from the body plan after the corrections. The manufacturability tests considered are *colliding*, *orientation* and *gripper access* tests. It is important to mention that the genome (CPPN) is not modified with this correction mechanism. The repair method only changes the genotype-phenotype mapping by repressing genes to be expressed in the final body plan.

4) Repair Method outside of the Evolutionary Loop (GRe):: In this algorithm the correction mechanism is applied to the population of robots generated with *NS*. Hence, this method does not affect the evolutionary process. The objective of this algorithm is to improve the final population with a post-fix mechanism.

IV. Results

This section provides answers to the three research questions (section 1), comparing algorithms in terms of the number of manufacturable robots generated and their diversity. All experiments use the same parameters: CPPNs are initialised with 10 hidden layers with 10 hidden neurons each. The total number of evaluations for each replicate (repetition) is 20,000 (500 generations with a population of 40) where the number of replicates is 15 for each experiment. Mann-Whitney u-test is used to demonstrate that two distributions are statistically significant different when $p < 0.05$.

A. Generating a population of manufacturable robots

The manufacturability property is a critical property of the evolved robots if the robots are going to be built in the physical world. Figure 2 shows the ratio of unique robots that fail each test for all algorithms, where a higher ratio indicates a higher number of robots that passed each test, hence more manufacturable.

		Robots	Percentage	(%)
	Alg	U	UM	UMV
1	RS	14551 ±982	6.8 ±0.4	2.2 ±0.2
2	NS	9827 ±311	11.3 ±2.6	1.8 ±0.3
3	GRi	9704 ±297	56.3 ±4.2	3.5 ±0.6
4	GRe	9917 ±469	52.1 ±3.4	2.2 ±0.6

TABLE I

This table shows the mean and standard deviation across the 15 replicates for each experiment. U is the number of unique robots, UM is the percentage of manufacturable robots and UMV is the percentage of manufacturable and viable robots. The highest percentages are highlighted.

The *RS* generates the highest number of robots with skeleton and the head organ connected to this skeleton. This is because the skeletons generated for these robots are large and similar to each other. There is no significant improvement across any of the *NS*, *GRi* and *GRe* experiments for the presence of skeleton and head organ connected to the skeleton. This is because these strategies do not address these types of failures, it can be seen from Figure 2 that this is not the case for collision, orientation and gripper access. *GRi* and *GRe* produced the fewest robots failing the collision, orientation and gripper access tests. This is due to the correction mechanism that removes any conflicting organ.

The percentage of 'unique' manufacturable robots for each experiment is shown in Table I. The 'unique' robots are all the robots that have different morphological descriptors within a run. The *GRi* and *GRe* algorithms generate the highest proportion of manufacturable robots. This demonstrates that the gene repression correction mechanism does lead to an increase in the number of manufacturable robots generated. The remaining non-manufacturable robots failed the *skeleton presence* and *skeleton connected to head organ* tests. Also, in-evolution repair has a higher number than post-evolution repair because the repair operation is done before evaluation and selection take place. This influences the evolutionary process, which is not the case with post-evolution. A remarkable outcome is the big difference between manufacturability and viability. Using a repair mechanism the percentage of manufacturable robots is above 50% (which we found surprisingly high), but the the percentage of manufacturable and viable robots is only 2 - 3 percent (which we found surprisingly low).

This section has illustrated that having a correction mechanism helps to increase the number of manufacturable robots. However, this does not guarantee that the evolved robots will be diverse w.r.t. their morphological descriptor. In the next section, how the diversity changes with each algorithm is explored.

B. Generating a population of diverse robots

Diversity is an important property for any population of robots that acts as a seed for a subsequent evolutionary loop. We measure diversity using the *sparseness* metric, calculated per robot as the average distance of its morphological descriptor to its $k = 15$ closest neighbours (for all unique robots). For each experiment, we then calculate the mean sparseness over all unique robots.

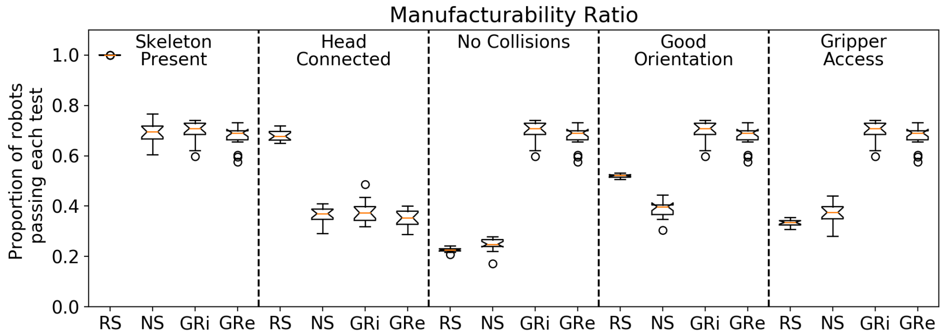

Fig. 2. The higher the proportion, the more robots have passed each test making them more manufacturable. The gene repression correction mechanism decreases the number of robots failing the no collisions, good orientation and gripper access tests. Boxplots follow the Tuke's original convention.

Sparseness ($\times 10^{-3}$)

	Alg	U	UM	UMV
1	RS	3.3 ±0.3	9.7 ±0.5	10.8 ±0.7
2	NS	6.7 ±0.4	11.7 ±0.9	18.8 ±2.2
3	GRi	6.0 ±0.5	6.2 ±0.7	13.7 ±1.7
4	GRe	5.8 ±0.3	6.1 ±0.4	17.3 ±2.2

TABLE II

THIS TABLE SHOWS THE MEAN AND STANDARD DEVIATION ACROSS THE 15 REPLICATES FOR EACH EXPERIMENT WHERE U REPRESENTS ALL THE UNIQUE ROBOTS, UM ONLY MANUFACTURABLE ROBOTS AND UMV ONLY MANUFACTURABLE AND VIABLE ROBOTS. THE HIGHEST SPARSENESS VALUES ARE HIGHLIGHTED.

NS outperforms the other methods in generating the most sparse robots (Table II). Since there are no manufacturability constraints, evolution can traverse the landscape more easily. It is more difficult for *GRi* to traverse the exploration space since the population becomes stagnated with basic manufacturable robots with a low number of organs.

The exploration space for the *width*, *depth* and *height* attributes of the morphological descriptor is mostly covered for all the experiments (Figures 3a, 3c, 3e and 3g). This is because simple configurations of the CPPN can generate robots of different sizes. This is not the case for the *wheels* and *sensors* attributes (Figures 3b, 3d, 3f and 3h). More complex configurations of CPPNs are required to combine multiple regions in the decoding stage to generate a high number of organs. In conclusion, the *NS* algorithm generates the greatest diversity of robots. The *wheels* and *sensors* exploration space is more difficult to traverse.

C. The diversity-manufacturability trade-off

In the previous section, it was shown that it is possible to generate the most diverse population of robots with *NS* and the largest population of manufacturable robots with *GRi*. In this section, the trade-off between diversity and manufacturability across each algorithm is explored in more detail. Each algorithm explores a different region in the manufacturability/ novelty space, as illustrated in Figure 4. *NS* and *GRi* generate the largest number of robots at different ends of the spectrum. *NS* generates the most sparse robots and *GRi* generates the highest number of manufacturable robots. Figure 5 visualises the 6-dimensional space using a dimensionality reduction technique t-SNE [19]. The first row illustrates all unique descriptors, regardless of manufacturability or density. The density reduces dramatically after all the non-manufacturable robots are removed (2nd row). Note that robots that are manufacturable are not necessarily functional. For instance, a robot with no components (wheels or sensors) will pass all the manufacturability tests but will not be viable. After removing all non-viable robots (robots with no sensors and wheels) the density of evaluations reduces even further as shown in the final row in Figure 5. *NS* has the highest sparseness of all the algorithms, while *GRi* produces the highest number of manufacturable robots. Examples of manufacturable and viable robots are shown in Figure 6.

V. CONCLUSIONS

One of the main goals of the ARE project is to accomplish autonomous robot evolution. Automated manufacture of diverse robots is a key feature to this end, but this introduces constraints on the body plans that can be produced [5].

In this paper, three different algorithms were evaluated in terms of their ability to produce diverse, manufacturable and viable robots. The results showed that unsurprisingly, the methods augmented with repair processes produce the highest number of manufacturable robots, while pure novelty search the most diverse, with a trade-off clearly apparent between the two. However, for the first time we have provided evidence to quantify these effects.

We find that the repair mechanism leads to approximately 56% of individuals being manufacturable in the best case, while only 3.5% of these individuals are viable, indicating that only a very small region of the search-space leads to robots of interest. Regardless of this, at minimum we are able to generate a diverse population of 175 manufacturable and viable robots, which is consistent with typical population sizes used in evolutionary robotics, therefore suggesting that this population could be used in future to seed a second round of evolution in which body plans and controllers evolve

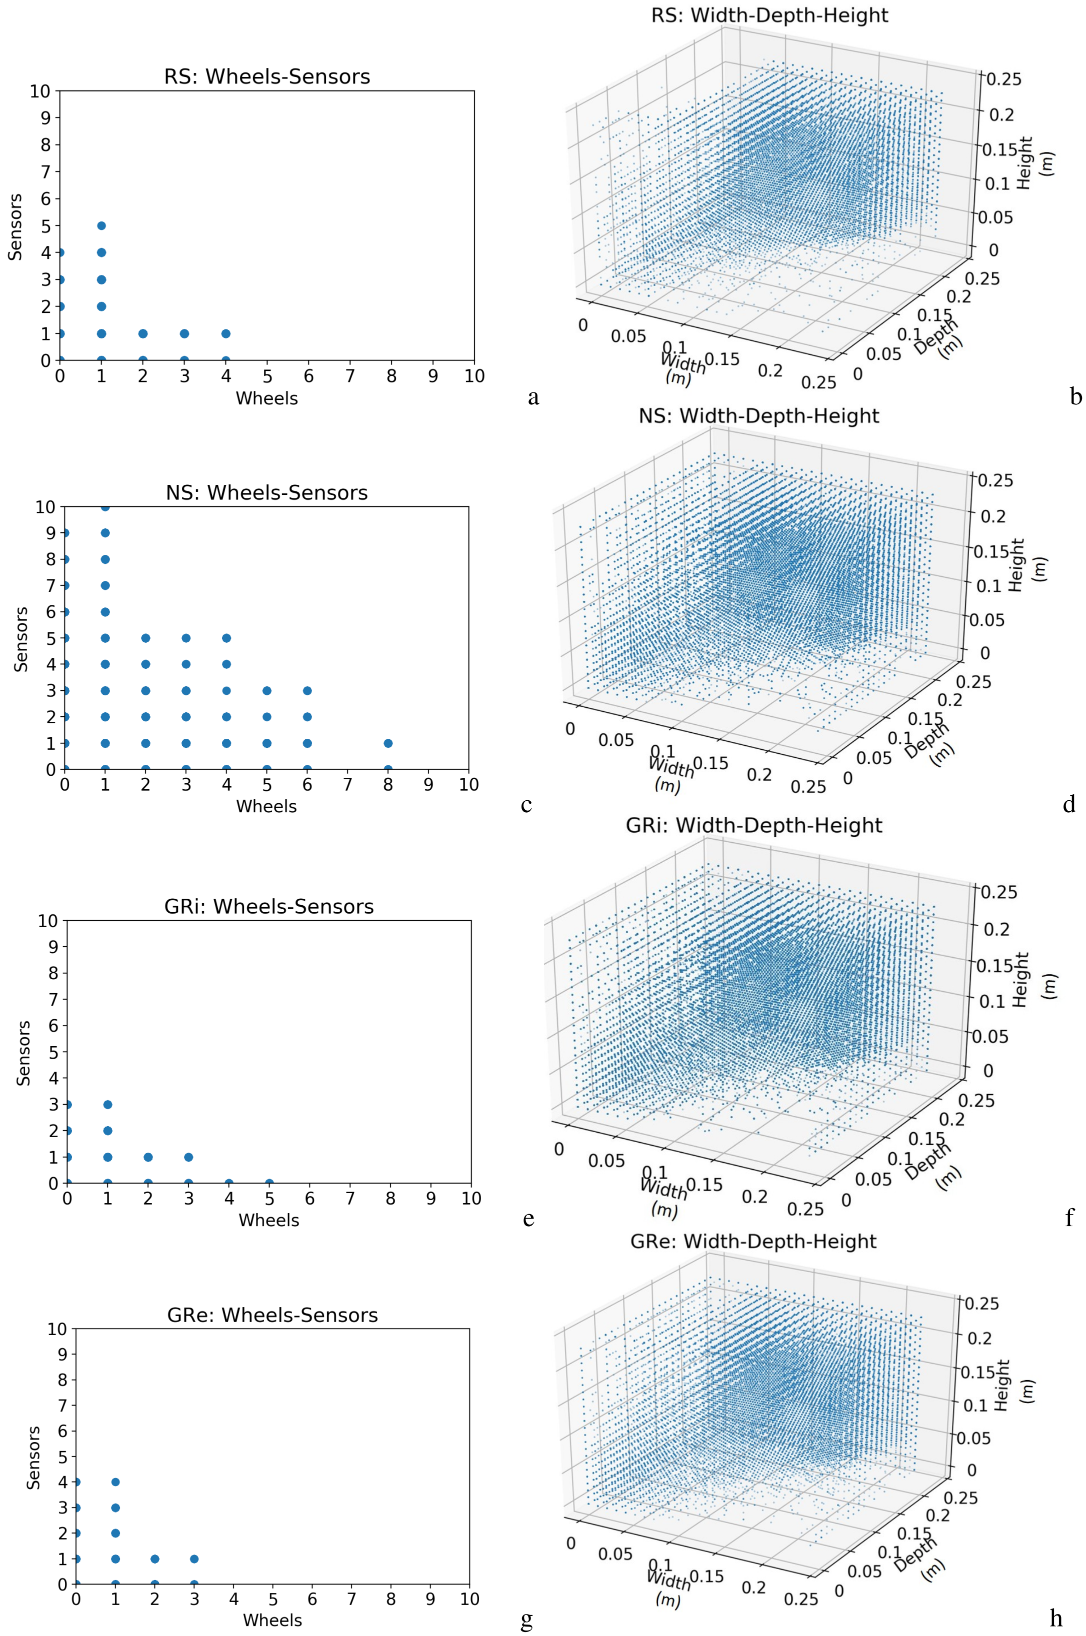

Fig. 3. Scatter plots for each different experiment (rows) and for different number of attributes of the morphological descriptor (columns). Each marker represents a robot. The exploration space is mostly covered for the *width*, *depth* and *height* attributes. This does not apply with the *wheels* and *sensors* attributes showing that it is more difficult to evolve robots with different number of organs. *NS* covers more space than the other algorithms.

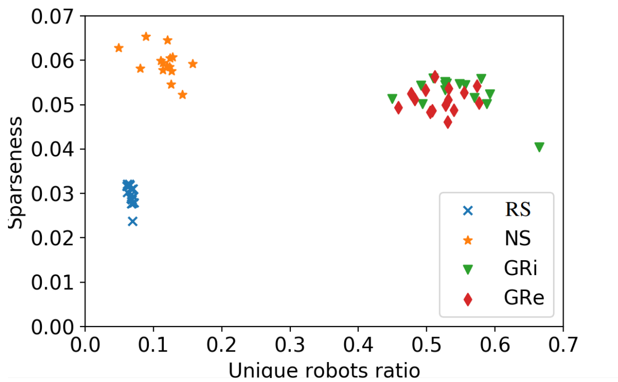

Fig. 4. Manufacturability/diversity trade-off. *NS* generates the most diverse robots and *GRi* the highest number of manufacturable robots.

together. With this, quite possibly, this population could save computation time for the second round of evolution.

As future work the body plan encoding will be improved to increase the number of manufacturable and viable robots. For instance, a variation of the encoding could make sure that there is always skeleton present and that is attached to the head organ. In addition, future work based on this study is concerned with using novelty search and repair mechanisms in an evolutionary process where selection is based on the behaviour of robots, not only morphological properties.

References

[1] A. Eiben, S. Kernbach, and E. Haasdijk, "Embodied artificial evolution – artificial evolutionary systems in the 21st century," *Evolutionary Intelligence*, vol. 5, no. 4, pp. 261–272, 2012.

[2] A. E. Eiben and J. Smith, "From evolutionary computation to the evolution of things," *Nature*, vol. 521, no. 7553, pp. 476–482, 2015.

[3] 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.

[4] D. Howard, A. E. Eiben, D. F. Kennedy, J.-B. Mouret, P. Valencia, and D. Winkler, "Evolving embodied intelligence from materials to machines," *Nature Machine Intelligence*, vol. 1, no. 1, pp. 12–19, 2019.

[5] 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 *The 2018 Conference on Artificial Life: A Hybrid of the European Conference on Artificial Life (ECAL) and the International Conference on the Synthesis and Simulation of Living Systems (ALIFE)*. MIT Press, 2019, pp. 95–102.

[6] 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.

[7] J. Lehman and K. O. Stanley, "Abandoning objectives: Evolution through the search for novelty alone," *Evolutionary computation*, vol. 19, no. 2, pp. 189–223, 2011.

[8] H. Lipson and J. B. Pollack, "Automatic design and manufacture of robotic lifeforms," *Nature*, vol. 406, no. 6799, pp. 974–978, 2000.

[9] 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.

[10] 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.

[11] S. Kriegman, D. Blackiston, M. Levin, and J. Bongard, "A scalable pipeline for designing reconfigurable organisms," *Proceedings of the National Academy of Sciences*, 2020.

[12] P. Funes and J. Pollack, "Computer evolution of buildable objects," *Evolutionary design by computers*, vol. 1, pp. 387–403, 1999.

[13] J. Collins, W. Geles, D. Howard, and F. Maire, "Towards the targeted environment-specific evolution of robot components," in *Proceedings of the Genetic and Evolutionary Computation Conference*, 2018, pp. 61–68.

[14] J. Hiller and H. Lipson, "Automatic design and manufacture of soft robots," *IEEE Transactions on Robotics*, vol. 28, no. 2, pp. 457–466, 2011.

[15] M. F. Hale, M. Angus, E. Buchanan, W. Li, R. Woolley, L. K. Le Goff, M. de Carlo, J. Timmis, A. F. Winfield, E. Hart, A. E. Eiben, and A. M. Tyrrell, "Hardware design for autonomous robot evolution," in preparation.

[16] K. O. Stanley, "Compositional pattern producing networks: A novel abstraction of development," *Genetic programming and evolvable machines*, vol. 8, no. 2, pp. 131–162, 2007.

[17] K. O. Stanley, D. B. D'Ambrosio, and J. Gauci, "A hypercube-based encoding for evolving large-scale neural networks," *Artificial life*, vol. 15, no. 2, pp. 185–212, 2009.

[18] J. Gomes, P. Mariano, and A. L. Christensen, "Devising effective novelty search algorithms: A comprehensive empirical study," in *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, 2015, pp. 943–950.

[19] L. v. d. Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of machine learning research*, vol. 9, no. Nov, pp. 2579–2605, 2008.

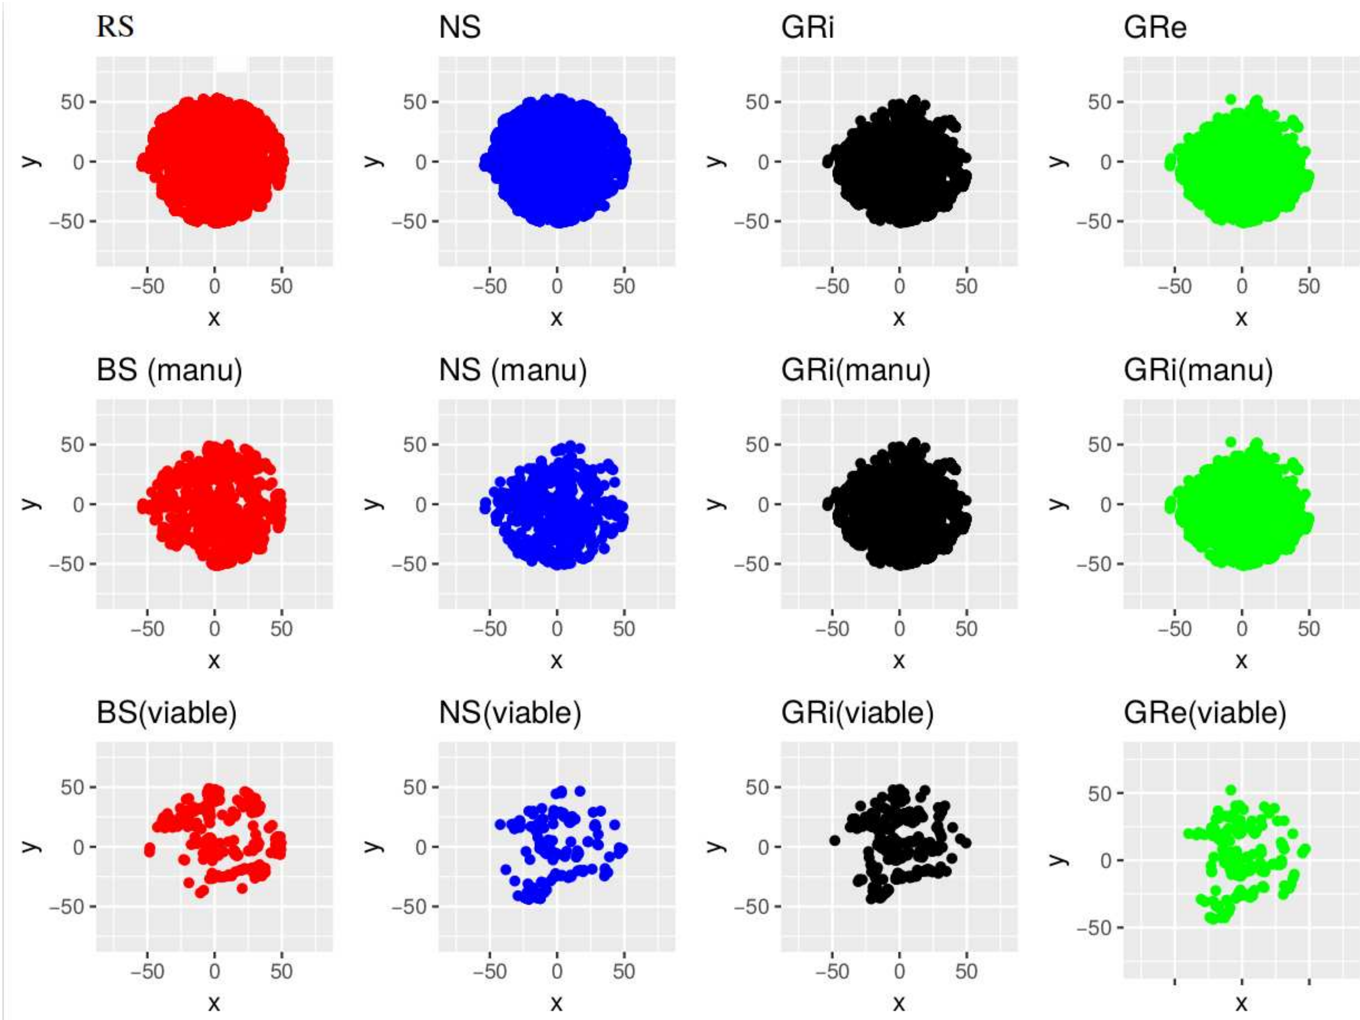

Fig. 5. Low-dimensional projection of the 6D descriptors: The first row shows unique robots, the second manufacturable robots and the last manufacturable and viable robots. The body plan density decreases with when only the manfuacturable and viable robots are considered. Also, it is possible to spot empty regions where all the non-manfuacturable robots exist.

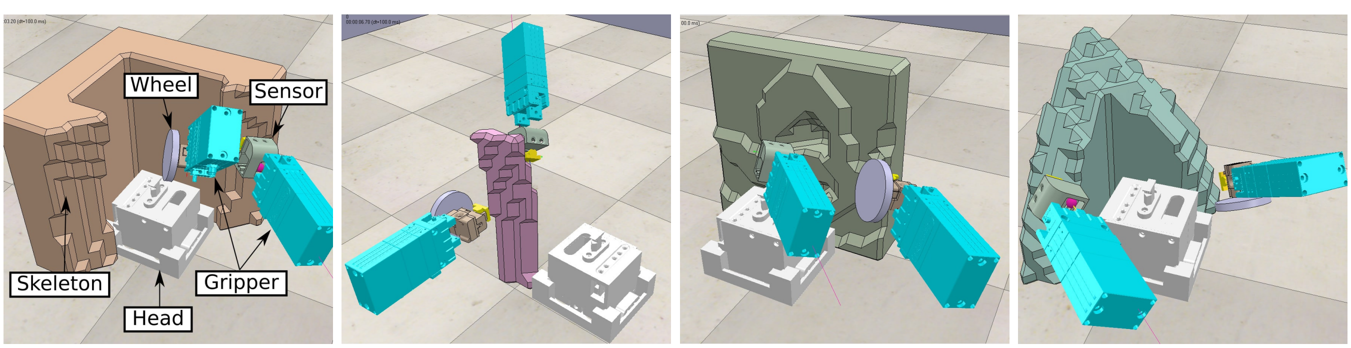

Fig. 6. Manufacturable and viable robots produced with each algorithm: *RS*, *NS*, *GRi* and *GRe* (shown in this order) Each robot has the skeleton connected to the head organ, clear access for the grippers to connect the organs and the organs are not colliding and have a good orientation. The grippers (blue components) represent the gripper of the robot arm in the Robot Fabricator assembling the robot