RESEARCH PAPER

Heritability of morphological and behavioural traits in evolving robots

Matteo De Carlo[1] · Eliseo Ferrante[1,2] · Daan Zeeuwe[1] · Jacintha Ellers[1] · A. E. Eiben[1]

Abstract

In the field of evolutionary robotics, choosing the correct genetic representation is a complicated and delicate matter, especially when robots evolve behaviour and morphology at the same time. One principal problem is the lack of methods or tools to investigate and compare representations. In this paper we introduce and evaluate such a tool based on the biological notion of heritability. Heritability captures the proportion of phenotypic variation caused by genotypic variation and is often used to better understand the transmissibility of traits in real biological systems. As a proof of concept, we compare the heritability of various robot traits in two systems, one using a direct (tree based) representation and one using an indirect (grammar based) representation. We measure changes in heritability during the course of evolution and investigate how direct and indirect representation can be biased towards more exploration or exploitation throughout the course of evolution. The empirical study shows that heritability can be a useful tool to analyze different representations without running complete evolutionary processes using them.

Keywords Evolutionary robotics · Morphological evolution · Representation · Heritability

1 Introduction

Evolutionary Robotics [1–3] is a research field in Artificial Intelligence that applies Evolutionary Computing methods to an embodied robotic platform. The main objective of the field is to use Evolutionary Computing to design robots that are ideal to specific tasks or environments. But, Evolutionary Robotics can also be used as an evolutionary model with physical plausibility given its ability to generate individuals with agency in a complex world. Because of this, we can use it to fill some gaps between simple evolutionary models and the enormous complexity of natural evolution.

The modular robots we use are defined by a genetic representation that not only defines the behaviour of the robot, but also its morphology [4]. In this paper we compare two different typologies of genetic representation: a simple one, which we will refer to as *direct representation*, and a complex one, which will be referred to as *indirect* or *developmental representation*, where there are rich interactions between genes. In our developmental representation, each phenotypic trait is the result of a complex series of interactions between different genes. The implementation details of the two representations will be further explained in the Methodology (Sect. 3). What is important to note is that, in our model, phenotypic traits are not transmitted directly between generations, but they are an expression of the underlying genotype of the individuals. The genotype is transmitted between generations through inheritance of the corresponding genes. This is especially relevant for the developmental representation, where every phenotype is an expression of an interaction of multiple genes. To compare these two different setups, we are interested in the investigation of *heritability* [5, 6], a property that is usually ignored in evolutionary models but often measured in real populations when studying the transmissability of phenotypic traits from parents to offspring.

✉ Matteo De Carlo
m.decarlo@vu.nl

Eliseo Ferrante
e.ferrante@vu.nl

Daan Zeeuwe
daan_zeeuwe@hotmail.com

Jacintha Ellers
j.ellers@vu.nl

A. E. Eiben
a.e.eiben@vu.nl

1   Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

2   Autonomous Robotics Research Centre, Technology
     Innovation Institute, Abu Dhabi, United Arab Emirates

We present this study as a proof of concept for heritability as a useful tool, to aid the design of Evolutionary Robotic systems. When designing an Evolutionary Robotic system, one important design choice that has a major impact but is often overlooked is the genetic representation. There are several representations available to the designer that have been used in the literature, and many more if one includes all variations. Nevertheless, our understanding of how the choice of a particular representation can influence the evolutionary process is still very superficial. By comparing two substantially different representations, we hope to investigate the applicability of the notion of heritability in an evolutionary robotics system to better understand the relationship between different representations and the generated phenotypes.

The main contribution is the adoption of the concept of heritability to Evolutionary Robotics and the demonstration of its utility in complex evolutionary models. There are four important aspects we investigate:

- Whether heritability can be used as a predictor of the evolutionary response of a system, specifically whether it is related to the rate of evolutionary change.
- Whether heritability changes over the course of evolution and, if so, if this effect can be related to any other measurable aspect of evolution.
- Whether heritability can be used to evaluate and compare different genetic representations.
- Whether heritability can help us better understand the different effects that simple and more complicated genotypic representations have on evolution.

2 Related work

Evolution is the process that produced complex life as we know it, and ultimately human conscience, but it took millions of years to arrive at the complexities that we know today. To study evolution, models are generally a good solution because they allow for simulation of evolutionary processes across multiple generations in a fraction of the time that real-life evolution takes place. In the past decades, many important insights have been derived from evolutionary models. For example, speciation models have helped identify the conditions under which sympatric speciation can occur, and how learning and phenotypic plasticity promotes or inhibits population divergence [7–9]. Evolutionary models have also greatly increased our understanding of sexual selection and the importance of condition dependence of sexual selected traits [10–12].

One aspect of evolutionary modelling that has been notoriously difficult, is the genetic representation of the traits under selection. Often, models simulate trait variation as a normal distribution with genetically controlled parameters [13,

14], or assume multiple genes with additive effects to yield a direct relationship between genetic parameters and the trait under selection [15]. In living organisms, however, phenotypic variation in traits often arises as an emergent effect from complex molecular and developmental interactions, and cannot be traced back to a single gene [16]. Particularly for quantitative traits, such as fecundity, locomotive behaviour, and body size, phenotypic variation results from gene-gene and gene-environment interactions. Complex genotype-phenotype relationships can significantly alter the trajectory and outcome of evolution because of the reduced importance of additive genetic variation in such systems [17]. Therefore we need more insight in how encoding a more complex genetic system can affect the properties of evolutionary models.

In this study, we modelled evolution with physically plausible robots using an encoding-decoding process that allows interactions of genes by design, a method that is the basis of a field called Evolutionary Robotics. Studying biology using a robotic model is not a novel idea [18–22]; In [19] when explaining his use of robots in biological research, the author argues that studying biological systems through the use of a model can be appropriate, when the element that we want to study is modelled properly.

Similar attempts of monitoring the variance in the phenotypic space have been proposed before. In relevant studies [23–25] *evolvability* is defined as "the propensity of an individual to generate phenotypic diversity", and is monitored and used to improve the results of evolution. In these studies the objective is to improve the efficacy of the evolutionary process by monitoring and improving the evolvability of the individuals. In [24] individuals with better evolvability improve the exploration of the system in conjunction with Novelty Search. In respect of our work, evolvability seem to focus on measuring the overall individual effectiveness on creating new interesting solutions, while heritability is the analysis of single traits and the understanding on whether those traits are transmittable, with distinct measurements for each trait. To summarize, evolvability is an estimation on how frequently new interesting mutations appear, while heritability measures how new helpful mutations are likely to survive multiple generations. Therefore, while related, evolvability and heritability are two sides of the same coin, and should be valued when designing an evolutionary process.

3 Methodology

3.1 The robots

The platform we used is a modular robotic framework based on RoboGen [26]. Each robot is composed of three different types of modules: one Core module (Fig. 1a), an arbitrary number of Brick modules (Fig. 1b), and an arbitrary number

of Joint modules (Fig. 1c). The Core module is unique for each robot and represents the robot "head" that, in the original physical incarnation [27], contains the main logic board and the battery. The Core module has four connection points where other modules can be attached. Brick modules represent the "backbone" of the robot. Only through Brick modules, the robot can take up arbitrary shapes. Actuation can only be achieved through the Joint modules, thus Joint modules are the only modules capable of changing the state of the robot in the environment. Joint and Brick modules can be attached to any other module in two different ways, which differ from each other by 90° for the axis perpendicular to the attachment plane. In [28], we already introduced this rotational attachment, but it applied only to Joint modules. Allowing the Joint to be attached rotated permits the robot to evolve morphologies that have more variety in terms of both degrees of freedom and actuation. In this paper, we also introduce rotational attachment for Brick modules, which potentially allows robots to extend vertically against gravity. In general, the design allows the inclusion of sensors, but for this study, we do not use any. An example of a four-legged spider robot is presented in Fig. 2a.
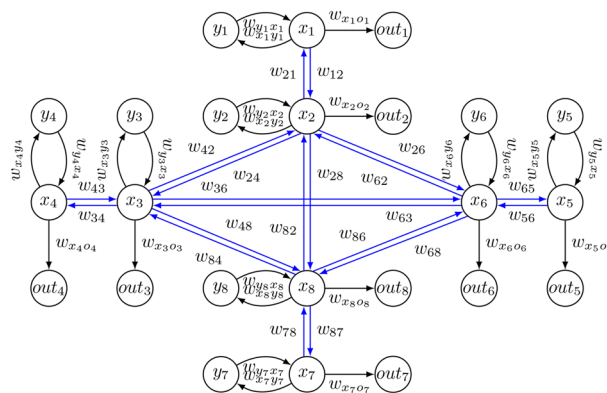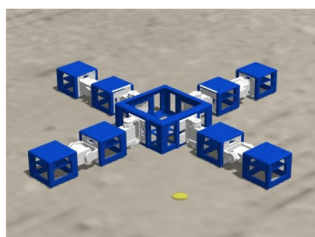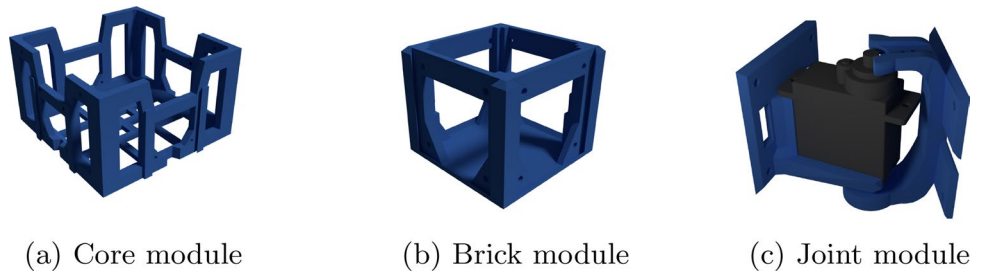
3.2 Robot brains

The controller of the robots we used is based on Central Pattern Generators (CPG) [29–32]. The CPG implementation we used in our work is from Lan et al. [33, 34]. This architecture has proven to be quite successful with joint based robots; in our previous work [35], we successfully used this brain architecture to evolve two very different gaits: locomotion and rotation.

In our CPG network implementation, every joint in the body has a corresponding CPG node that consists of three neurons. Two of these neurons (that we call x and y neurons) are coupled by two-directional connections, one from x to y, and one from y to x. By definition, the weights of these connections have the same value but the opposite sign. The remaining neuron in a CPG node provides the output signal to the servo motor driving the given joint. The corresponding weight is set at 1.0 in each joint, thus a CPG node can be configured by just one parameter regulating the connections between x and y.

The overall controller architecture is a network with one CPG node for each joint and a connection between two such

Fig. 1 Robots are built using three types of modules. Starting with only one Core module, the robot grows by connecting Brick and Joint modules to the Core or other already connected modules

(a) Core module   (b) Brick module   (c) Joint module

(a) Spider robot example   (b) Spider robot brain diagram

Fig. 2 Example "spider" robot morphology with eight joints and corresponding brain diagram. In the brain diagram (2b), black arrows represent weighted connection between neurons in the same oscillator, while blue arrows represent weighted connections between neu-

rons of neighbouring oscillators. All connections between oscillators, inside and outside, are corresponded with a symmetrical connection of the same weight but opposite sign. (e.g. $w_{x_1 y_1} = -w_{y_1 x_1}$) (colour figure online)

nodes if the corresponding joints are neighbors separated by no more than two empty cells (in the Manhattan sense) in the 3D Euclidean grid enclosing the robot. Connected neighbor CPG nodes can synchronize the oscillations of their joints and induce global locomotion patterns. The number of configurable parameters for a robot brain is thus $j + c$, where j is the number of joints, and c is the number of connections between joints; in Fig. 2 we show an example of how the nodes would be connected in a robot made of 8 joints configured as in the "spider" robot from [27].

3.3 Evolution with the tree-based direct representation

For the direct representation, we are using a tree-based representation, whose implementation is very similar to [4, 36].

In a tree-based representation, the genotype is a tree data structure in which each node represents a module of the robot. Modules differentiate into three types: Core, Brick, and Joint modules, as represented in Fig. 1. The Core module is always the root of the tree and it can only be present once in the entire genotype. It can have four children. The Brick module is attached on one side to its parent block and has three remaining slots available for child nodes. The Joint module has only one remaining slot for a child module, therefore it does not allow any branching. A Joint has three extra parameters that directly encode the oscillator parameters: frequency, offset, and amplitude. In our tree-based representation, the brain development is limited to only decode parameters for the oscillators of the CPG network; i.e. all connections between oscillating nodes are not activated.

The robot module tree can be altered by either a body mutation or a brain mutation operator. Most mutation operations primarily revolve around changing the body: adding a random module, deleting a sub-tree, duplicating a sub-tree, or swapping a sub-tree. Brain mutations are achieved by mutating the joint oscillator parameters to achieve different activation patterns.

Parent robot trees can be recombined by inheriting subtrees from the parents. Some checks and balances ensure that the recombined trees are valid, and do not exceed the maximum limit of modules.

3.4 Evolution with the developmental representation

For the developmental representation, often called indirect representation, we choose a system from our previous work [28], which is composed of an Lindenmayer-system (L-system) [37] that describes body and brain structure and a component based on Hypercube-based Neuro Evolution of

Augmenting Topologies (HyperNEAT) [38, 39] that encodes the weights of the CPG controller network. These indirect representations are capable of creating symmetrical growth structures and repetitions in our bodies.

L-systems are parallel rewriting systems acting on a formal grammar. The grammar is defined as a tuple $G = (V, w, R)$, where V is the Alphabet, w is the Axiom and R is a set of Replacement Rules. L-systems start from the Axiom w, which is a sequence of symbols from the Alphabet. To develop an L-system grammar, the Axiom is expanded into a longer sentence by replacing symbols using the Replacement Rules in R. The replacement operation can be repeated multiple times on the sentence.

In this work, we adapted a system from [40], where each genotype is a grammar with always the same Axiom and Alphabet for all robots. The Alphabet is made of the following symbols:

- Robot modules: the Core, the Brick, a Vertical Joint, and a Horizontal Joint.
- Mounting commands: *add_left*, *add_front*, and *add_right*. Mounting commands must be followed by a module symbol otherwise they are ignored. When the final sentence is read, their role is to attach the following module symbol in the sentence to the module indicated by the cursor position, at a new position (left, front, or right) depending on the specific command.
- Moving commands: *move_back*, *move_right*, *move_front*, and *move_left*. The moving commands alter the position of the cursor.

The Replacement Rules of our L-system are a set of rules that replace any of the robot module symbols with a sequence of new symbols from the alphabet. In other words, robot module symbols are both terminal and non-terminal symbols in our L-system. Any other symbol is terminal, which means it cannot be replaced further. The Axiom of our L-system is a sentence made of a single symbol: the Core block. Once the L-system grammar develops the Axiom into a sentence, the sentence is used as a sequence of instructions that describe how to build the robot.

The system described in this paper differs from the one we inspired upon by the addition of an extra constraint on the morphology, i.e. we do not allow a joint to be attached to another joint. This new constraint increases the chances for more complex robots to appear, as shown in previous experiments [28]. Therefore, we decided to use it here so to allow the development of interesting morphologies in both experimental configurations. We also improved the Alphabet with the introduction of a new rotated Brick module, which allows for the morphologies to develop in three dimensions.

The brain structure is defined by the body; each joint creates a corresponding CPG oscillator node and connections are made using the rules previously explained (Sect. 3.2). When all connections are defined, each CPG node is positioned on a substrate space with x, y, z, w coordinates and each connection weight value is queried from a Compositional Pattern-Producing Network (CPPN) [41], as defined by the HyperNEAT algorithm. In the substrate space, x, y, and z determine the position of the CPG's node corresponding joint, while the w axis determines the front and back neurons in an oscillating CPG node.

The mutation and crossover operators are defined by their individual components: for the L-system component, we use the operators defined in [42]. For the HyperNEAT component, we use the operators as defined by HyperNEAT, with the exclusion of the species, i.e. genomes are not divided into species and crossover is possible for each pair of genome in the population.

3.5 Heritability

In biological systems, heritability distinguishes those phenotypic features that are an expression of genetic material from features that influenced solely by the environment. It's expressed as a value between 0 and 1, ranging from terrible to perfect heritability. In biological systems, estimated heritability for life history traits and behaviour are typically low to medium (ranging up to 0.30 [43]), whereas morphological traits are often found to have higher heritability (average heritability of 0.46 [44]). In an evolutionary system placed on a flat empty environment and designed with very simple gene encodings, where each element of the genome is the direct value of a trait, e.g. speed, size directly encoded as numbers, we expect perfect heritability for traits, with values of 1.0. On the opposite extreme, a trait with a low (0.1) or null value of heritability cannot be predicted from the parents and will be either the result of a complicated epistatic gene interaction, or the result of heavy environmental influences on the trait. Additive genes are those genes that have a direct impact on the phenotype and can be easily derived from the parents, e.g. if the height of a species is encoded with additive genes, the height of the offspring can be predicted from the average height of the parents. Other types of genes can have effects of dominance or epistasis on the phenotypes, i.e. a combination of different genes needs to align to have effect on the phenotype. The non-additive genes are much harder to predict and measure. Finally, a phenotype could be an expression of the environment and totally independent of the genes, e.g. a dog is more friendly if it grows up in a house. Heritability is commonly used to distinguish environmentally induced traits from inherited ones. In this work, we have an extremely simple environment, which causes negligible effects on the phenotype. Therefore we can use heritability to deduce the additive and non-additive components of robots' genes.

3.5.1 Formal definition

Following, we provide a detailed explanation of how heritability is defined and how it is measured. In an evolutionary system the phenotypic variation (V_P) of a population of individuals is an expression of genetic variation (V_G) and environmental factors (E).

$$V_P = V_G + E \tag{1}$$

The genetic variation can be further subdivided into three major components: additive genetic variation (V_A), non additive genetic variation caused by epistatic genes (V_{NA}) and effects of random mutations (M).

$$V_G = V_A + V_{NA} + M \tag{2}$$

As defined in [5, 6], heritability measures the contribution of genes to phenotypic traits. Each phenotypic trait has a different value of heritability. Heritability can be defined as broad-sense heritability (H^2) or narrow-sense heritability (h^2). Broad-sense heritability is the proportion of phenotypic variation that is created by the genetic variation. Narrow-sense heritability is only the proportion of genetic variation that is generated by additive genetic values, not including any effect of dominance or epistasis.

$$H^2 = \frac{V_G}{V_P} \quad h^2 = \frac{V_A}{V_P} \tag{3}$$

Heritability is also an important component of the "response to selection" (R), a value that can be predicted as the product of narrow-sense heritability and selection differential (S) [45]:

$$R = h^2 \cdot S \tag{4}$$

3.5.2 Measuring heritability

The value of heritability can be calculated from its theoretical formula, but this requires a deep mathematical understanding of our genotype model, which we do not posses. However, if we are only interested in the additive genetic material, an estimate of narrow-sense heritability can easily be derived from population measurements by linearly regressing the average trait value of the offspring against the parental phenotype. It is important that heritability is calculated between two consecutive generations and not more. An approximation to a linear model is possible because the additive genetic code has a linear response

to the resulting phenotype, in contrast to epistatic genetic code which has a much more unpredictable effect on the phenotype. The value for the slope of the linear regression is our numerical estimation for heritability. The value for heritability can vary between 1.0 and 0.0, where $h^2 = 1.0$ is a 45° linear regression, representing a perfect match between parents' average trait and the offspring's trait. A value of $h^2 = 0.0$ instead represents a scenario where the offspring's trait is completely unpredictable given the parents' traits.

In evolutionary robotics, the influence of the environment over the development of the individuals is usually very limited, excluding a few exceptions [46, 47]. In our case the environment is a flat terrain, therefore its influence is completely absent. The implication is that the phenotypic variation in this system is only an expression of genotypic variation. Through linear regression we can estimate the narrow-sense heritability, which is only an expression of additive genetic variation. The rest of the phenotypic variation can only be an expression of epistatic gene interaction and mutation (Eq. 2). The use of a linear regression is sufficient because we are mostly interested in an overall indication on how well certain phenotypic values transfer from parents to offspring.

3.6 Robot traits

To estimate an overall heritability of our system, we chose a wide variety of phenotypic traits that are representative of different aspects of our robot. This work is interested in the overall evolution of modular robots, but the approach is not limited to any particular number or type of traits. The same study can be repeated on any trait, e.g. it would be interesting to study the heritability of "the number of feet in a robot" and what parameters increase the transmission of the trait to the offspring.

We recorded a set of many traits derived from the descriptors found in [47]. From the many traits available, we sampled only a significant few that we found to be orthogonal to each other in previous work [48]: some traits that measure the morphological aspect of the robots and some that measure the behavioural aspect.

The **Morphological** traits give us insight on how the robot shapes evolve. In this work we used:

- *Proportion*: considering the 2D bounding box that encompasses the robot when viewed from above, this trait is the ratio between the two sides of this rectangle.
- *Size*: the number of modules in the body.
- *Number of Limbs*: considering the robot as a tree of modules, it is the number of leaf modules. The value is normalized per robot by the number of all possible limbs available.

- *Coverage*: considering the 3D bounding box that encompasses the robot, this trait is the ratio between the area that is occupied by modules and the total area of the rectangle.

The **Behavioural** traits are very important because they give insight in the complex relationship of body and brain. In this work we used:

- *Speed*: Describes the average robot speed (*cm/s*), and is calculated as if the robot took the shortest path from the start position s_0 to the end position s_t, and is defined with Eq. 5. In our setup, the average robot speed is also used as the fitness of the robot.

$$v_{disp} = \frac{s_T - s_{t_0}}{\Delta t} \tag{5}$$

- *Balance*: We use the rotation of the head in the *x–y* plane to define the balance of the robot. We describe the rotation of the robot with three dimensions: roll ϕ, pitch θ, and yaw ψ. Thus, we consider the pitch and roll of the robot head, expressed between 0° and 180° (because we are not interested in whether the rotation is clockwise or anti-clockwise). Perfect balance corresponds to $\theta = \phi = 0°$, so that the higher the balance, the less rotated the head. Formally, balance is defined by Eq. 6.

$$b = 1 - \frac{\sum_{t=1}^{T} |\phi_t| + |\theta_t|}{180 \cdot 2 \cdot \Delta t} \tag{6}$$

4 Experimental setup

We evolve the robots for movement, defining movement as how far robots travelled from the starting position, the direction being not important. We use the *speed* behavioral trait as the value for fitness, as shown in Eq. 5, which presents an easy challenge that we know it can be solved by our setup from previous experiments. For both representations we use the same evolutionary algorithm with a generational population update scheme, that is, an evolutionary algorithm where consecutive populations are non-overlapping. This means that survival selection is trivial: no members of population P_n survive, the subsequent generation P_{n+1} consists of offspring of the current one. As for parent selection, we use the tournament selection mechanism with a tournament size of two individuals. This represents a very low selection pressure. To generate more consistent data for heritability, every new individual is generated from the crossover of two parents, i.e. no individual is generated by mutation only. We run this algorithm with a population size of 100 individuals for

50 generations, amounting to a total of 5000 evaluations as the computational budget for optimizing the robots' makeup. 50 generations are not a lot in the evolutionary robotics field, and we do not expect evolution to converge, but it is enough for the type of analysis we want to make. For both representations, fitness evaluations are done by placing the given robot on a flat surface and running it for 30 s. As mentioned above, fitness is measured as the distance between the starting point and the end point of the life of the robot divided by the lifetime in seconds. The resulting fitness is a speed value measured in *cm/s* and it has been used successfully in previous work to evolve robot that have a meaningful locomotion gait.

The direction of the robot is not important. Evolutionary robotics has the tendency to evolve robots that are as big as the system allows, therefore, we fixed the maximum number of modules per robot to 50 to limit the bloat effect and to have reasonable simulation times. Additionally, a very big robot is impractical to build in hardware and would collapse on his own weight, therefore undesirable. Both configurations were replicated 10 times, for a total of 20 experimental runs and 100,000 different robots simulated. Compared to previous experiments, we adjusted the mutation rates for the evolutionary run to be quite high, with a probability of 0.59 of having at least a body mutation for the tree-based representation and we used the same probability (0.59) for the mutation chance on the L-system grammar. The increased mutation rates increases the chance to explore more areas of the genotype space, allowing us to collect more data for our heritability study. More details on the mutation parameters can be found in Table 1.

Table 2 Estimated values of heritability for each phenotypic trait

Trait	Tree-based heritability	L-system heritability
Speed	0.74	0.35
Balance	0.77	0.37
Proportion	0.65	0.41
Size	0.73	0.47
N. of limbs	0.83	0.67

5 Results

Our first analysis aims at measuring the heritability of various traits in the two representation schemes at the start of the evolution experiment. To do so, for each representation/trait pair, we measure heritability using data from all evolutionary runs, but only on the very first generation. Heritability is measured comparing the trait value of an offspring against the average of the parents, therefore we need the offspring from the second generation as well. A linear regression is applied to the trait values of parents against offspring, and the slope of the resulting linear model is our estimate of heritability. The estimated values from our measurements are reported in Table 2.

The scatter plots used for estimating heritability can be seen in Figs. 3, 4 and 5. As we can see, the Tree-based representation consistently shows higher values of heritability for all traits.

Table 1 Mutation probabilities and parameters for the Direct-tree representation and the L-system+CPPN indirect representation

L-system and CPPN mutation prob.		Direct-tree mutation prob.	
Mutate L-system grammar	0.59	Duplicate random sub-tree	0.2
Mutate CPPN	1.0	Delete random sub-tree	0.2
Mutate weights	0.9	Generate random sub-tree	0.2
Add link	0.08	Swap two random sub-trees	0.2
Add neuron	0.01	(Derived) probability of at least one body mutation	0.59
Remove link	0	Mutate brain	0.2
Remove neuron	0	Mutate oscillator	0.5
Make recurrent link	0		
Make recurrent loop	0		
CPPN mutation params		**Direct-tree mutation params**	
Max weight	8.0	Mutate oscillator amplitude σ	0.3
Weight mutation max power	0.2	Mutate oscillator period σ	0.3
Weight replacement max power	1.0	Mutate oscillator phase σ	0.3

Values are derived from past experiments which had reasonable results. Mutation rates were increased substantially to increase the exploration and therefore collect more data points for our heritability analysis. In addition, the L-system grammar mutation was set to have the same probability of having at least one mutation in the Direct-tree representation
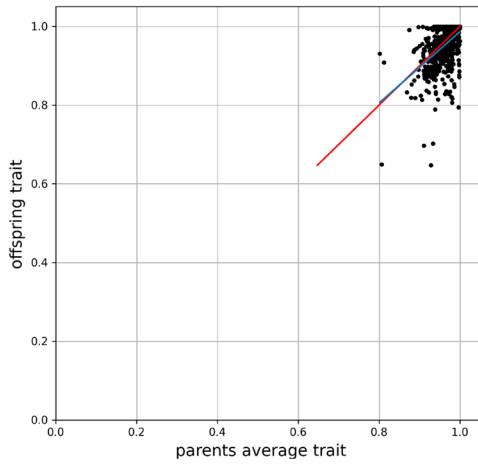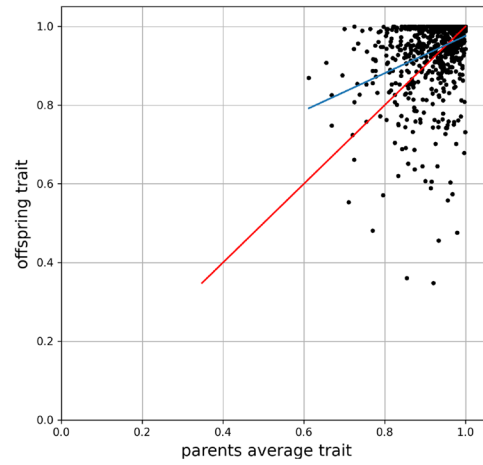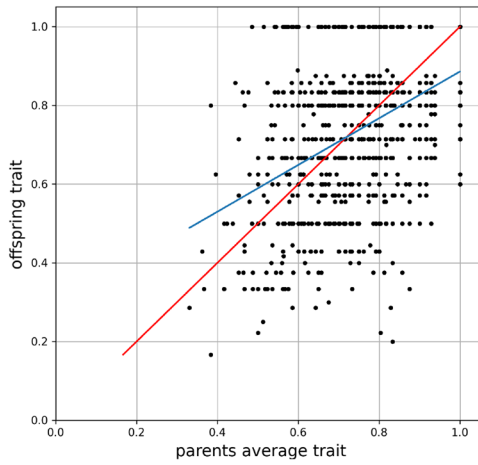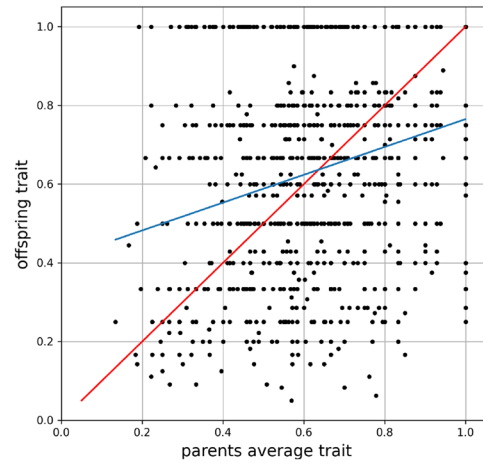
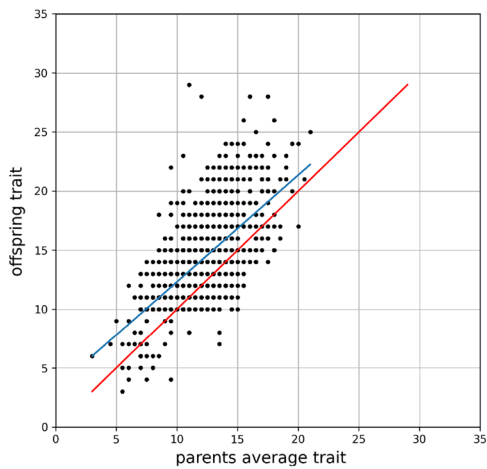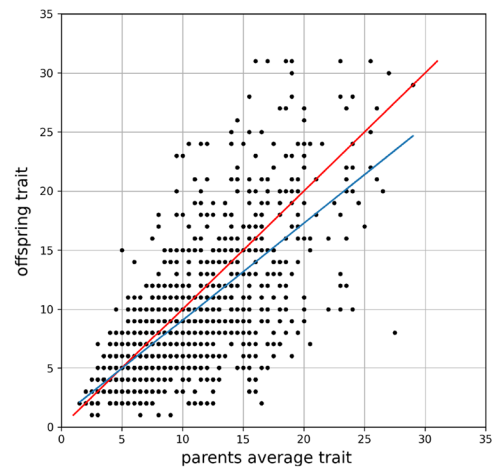(a) Balance: Tree–based

(b) Balance: L-system

(c) Proportion: Tree–based


(d) Proportion: L-system


(e) Size: Tree–based

(f) Size: L-system

◄**Fig. 3** Scatterplot for the two representations for the parents in the first generation and their offspring. The x axis shows the average trait value of the parents, the y axis shows the value of the offspring. The blue line represents the linear regression of these values. A steeper line indicates a higher level of heritability, to a theoretical maximum of 1 (45° slope). The red line is a 45° reference line (colour figure online)
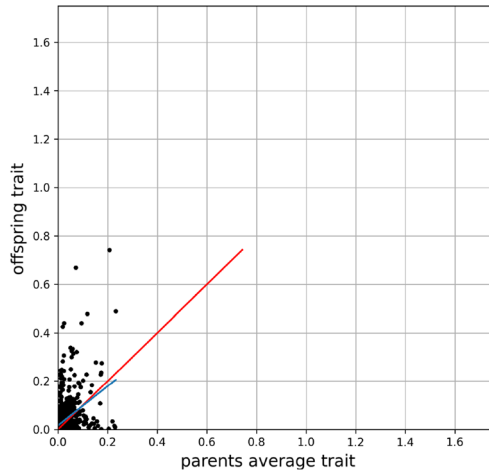
5.1 Relationship between heritability and initial evolutionary response

In this section, we aim at analyzing the relationship between heritability and evolutionary response for the

(a) Heritability of tree–based representation for the speed trait. $h^2 = 0.74$.
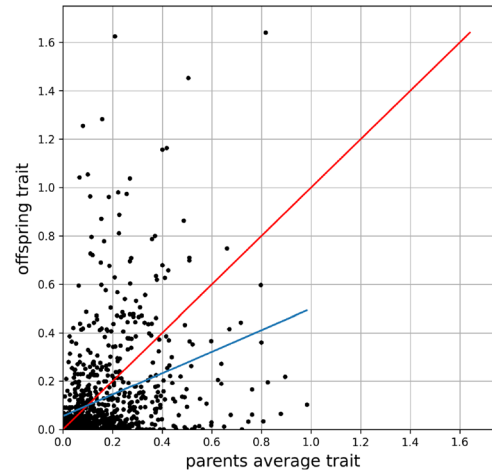
(b) Heritability of L-System representation for the speed trait. $h^2 = 0.35$.

Fig. 4 Heritability of the speed trait for both encodings measured at the first generation. Blue line is the resulting linear regression (slope is heritability). Red line is a perfect heritability reference. We can

observe that the data points are more scattered and that a lower heritability is measured in the L-system representation (colour figure online)
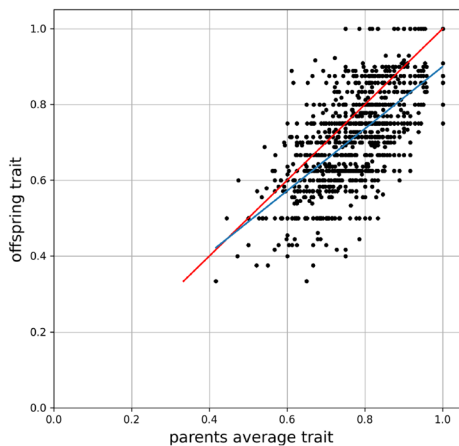
(a) Heritability of tree–based representation for the "number of limbs". $h^2 = 0.83$.
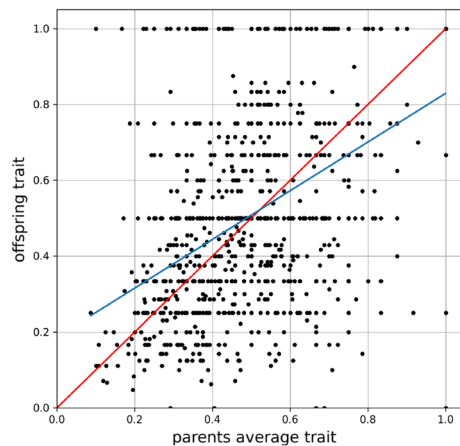
(b) Heritability of L-System representation for the "number of limbs". $h^2 = 0.67$.

Fig. 5 Heritability of the "number of limbs" trait for both encodings measured at the first generation. Blue line is the resulting linear regression (slope is heritability). Red line is a perfect heritability reference. We can observe how in the tree-based representation (5a)

points are well clustered around the 45° line, resulting in a pretty high value of heritability. In the L-system based representation (5b) points are instead scattered all around the plot, resulting in a lower value of heritability (colour figure online)

two representation schemes. We present the heritability for direct and indirect representations, calculated only in the first generation on one behavioural (*speed*, Fig. 4) and one morphological trait (*number of limbs*, Fig. 5). In Figs. 6 and 7 we show the dynamics of the two same traits. In both Figures, panel (a) shows the value of the trait over generations while Panel (c) shows its rate of change (or derivative) over generations. To further understand how heritability can explain the trait dynamics over generations, we use Panels (b) to highlight the evolution over generations of the heritability metric (calculated across two consecutive generations) and Panel (d) to highlight the evolution over a generation of the phenotypic diversity of the trait being considered within the population.

We first analyze the most important trait, since this is the one that is under selection: speed. In Fig. 6a, c, we observe that in the first 10 generations of the evolutionary process the Tree-based representation has a higher rate of change in fitness compared to the L-system. This finding is further confirmed by looking at the fitness distribution at Generation 0 for all runs and for the two representations, shown in Fig. 8.

Here we see that the L-system even starts with an advantage, represented by the much higher fitness diversity in the initial population and the corresponding presence of higher fitness individuals. Despite this, the L-system experiment evolves initially at a slower rate than the Tree-based experiment (as in Fig. 6c), which means that those high-fitness
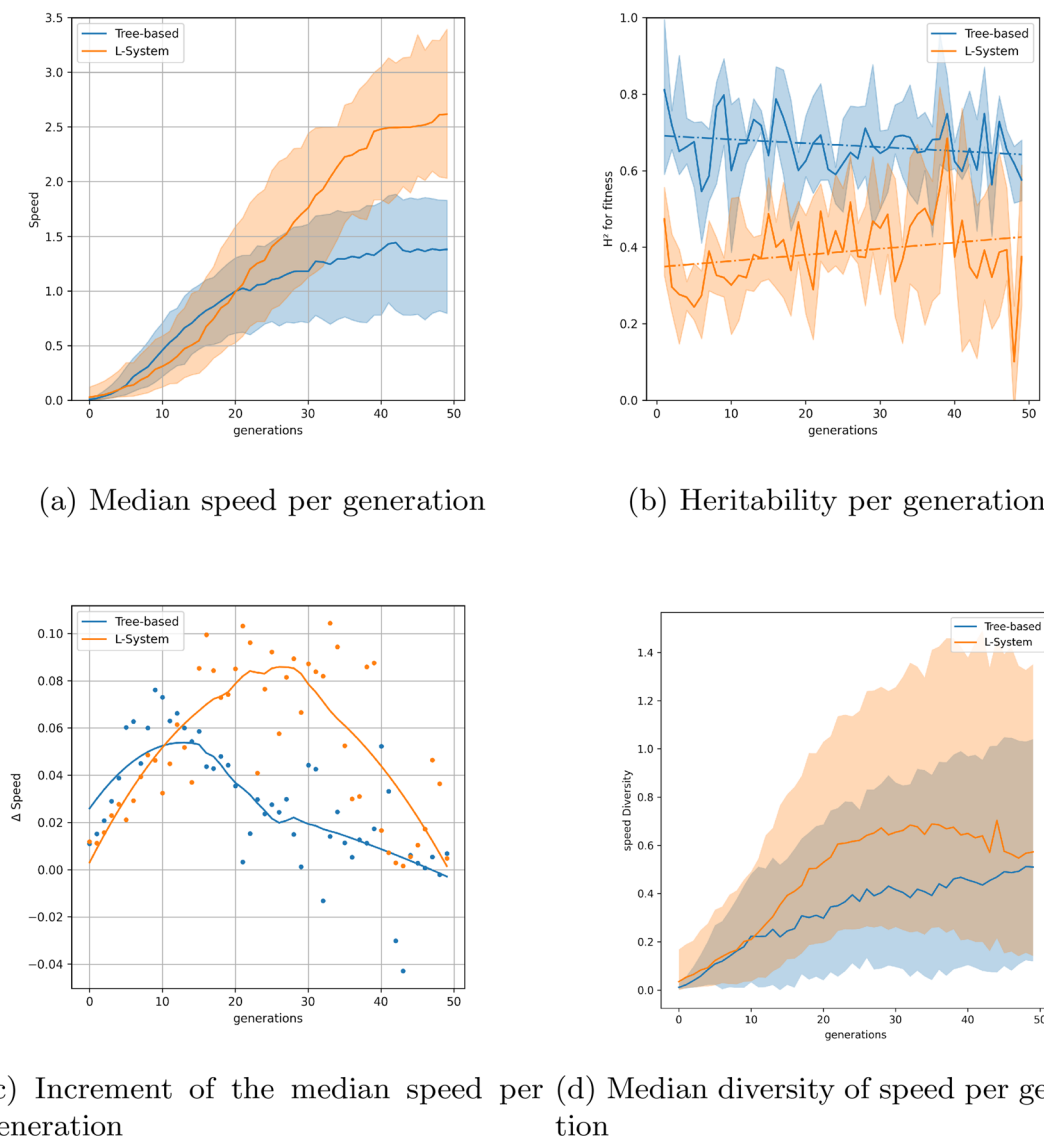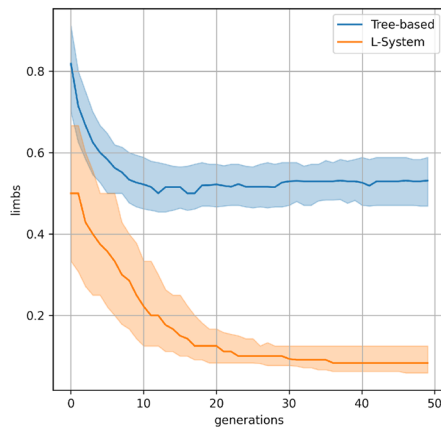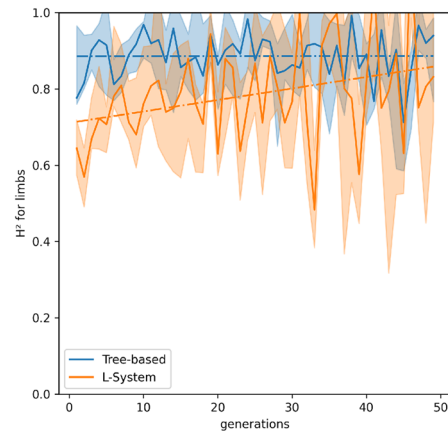

(a) Median speed per generation

(b) Heritability per generation

(c) Increment of the median speed per generation

(d) Median diversity of speed per generation

Fig. 6 Analysis of behaviour. Here we consider the evolution of the speed trait, how fast it changes per generation, speed heritability at the first generation and how it changes during evolution, and the diversity of speed present across the population
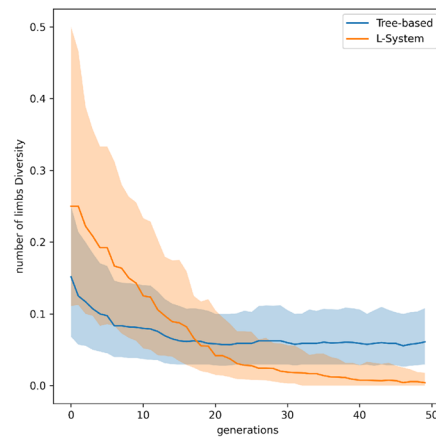
(a) Median number of limbs

(b) Heritability per generation
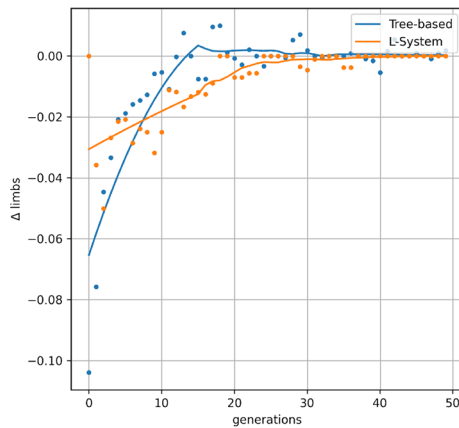
(c) Increment of the median number of limbs per generation

(d) Median diversity of number of limbs per generation

Fig. 7 Analysis of morphologies. Here we consider the evolution of the "number of limbs" trait, how fast it changes per generation, the trait's heritability at the first generation and how it changes during evolution, and the diversity of the "number of limbs" present across the population. The value for the "number of limbs" is normalized to the maximum possible number (of limbs) for each robot

individuals present in the populations are not able to pass their phenotype to their offspring to the same extent as it happens in the Tree-based representation. We argue that the concept of heritability can shed light to understand better what we observed above. Indeed, the Tree-based representation has a higher heritability value than the L-system representation (0.74 > 0.35, for all values see Table 2). Heritability can inform us on how much of the phenotypic trait variation will be passed on from parents to offspring, therefore high heritability at the beginning of the evolutionary process can predict a higher rate of change in the trait under selection, as is happening in our system. Thus, higher heritability at the beginning of the evolutionary process directly facilitates the effect of initial selection, because good parents have a higher probability of creating good offspring.

Importantly, the relation between initial heritability and the initial rate of change of a trait is not true only for traits that are specifically under selection. To support this claim, we perform the same analysis for all traits, and we report here only one example. Figure 7b shows the rate of change of a morphological trait that is not under selection (number of limbs). Also here, we observe the same overall pattern: the Tree-based representation has a higher initial rate of change in this trait, consistently with having higher heritability (Table 2).
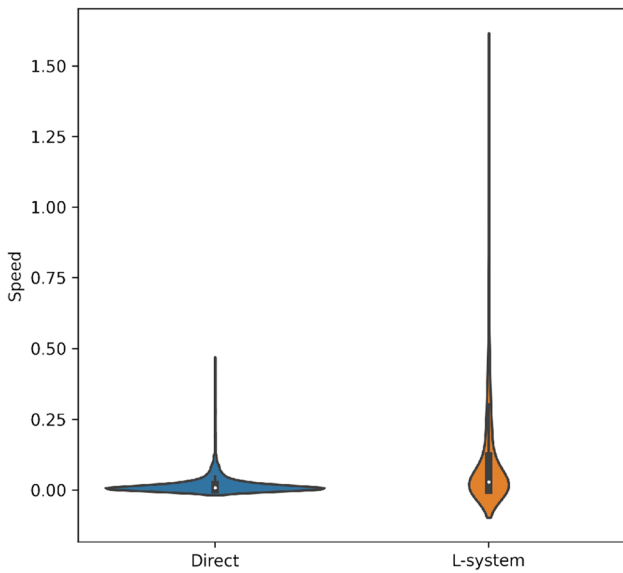
Fig. 8 Distribution of fitness in the random initial population

5.2 Heritability during the later phase of evolution

The analysis of heritability can not only help to describe the behavior of evolutionary systems in their initial phases, but it can also show interesting patterns during later phases of their process. From the theoretical definition, we expect the estimated value of heritability to be constant over the course of evolution, under the condition that the selection process does not affect the genetic variation in the population. Surprisingly, by computing the estimated value of heritability for each of the generations, shown in the top row of Fig. 9, we observe changes of the estimated heritability in our experiments. The change in heritability value across generations is most evident for traits that are not under selection. In the later stages of evolution heritability stabilizes for the tree-based representation and becomes highly unstable for the L-system genotype.
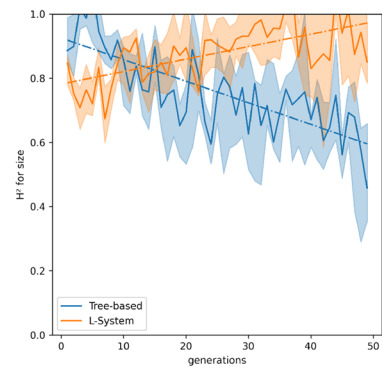
The change of heritability over generations can be explained if we also analyze how the phenotypic diversity
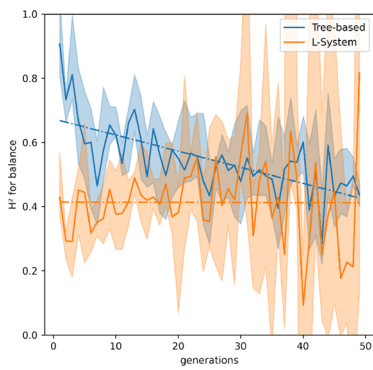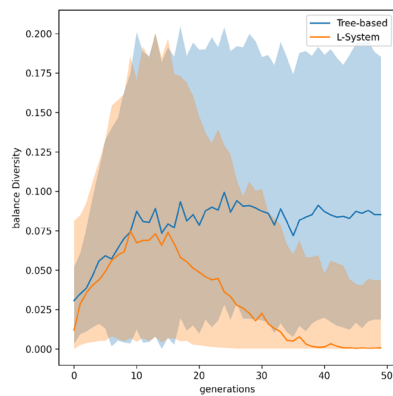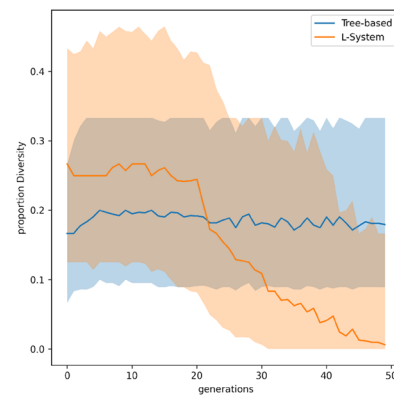

(a) Balance Heritability
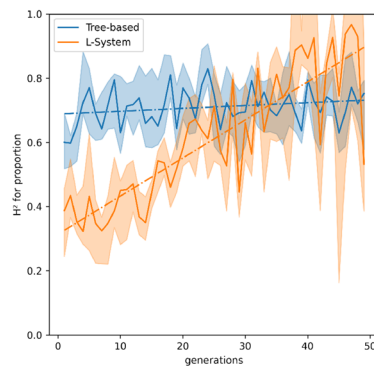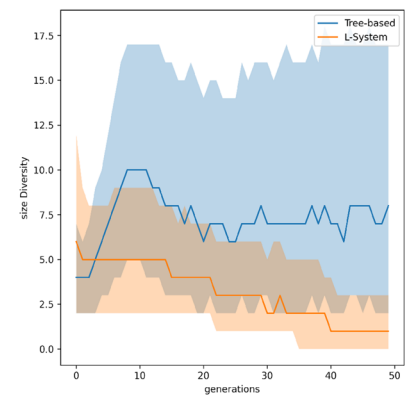
(b) Proportion Heritability

(c) Size Heritability

(d) Median diversity of balance per generation

(e) Median diversity of proportion per generation

(f) Median diversity of size per generation

Fig. 9 Heritability and diversity plots over generations for three different phenotypic traits: balance, proportion and size

Fig. 10 Diversity per generation. Diversity is the average distance of one individual against all other individuals in the same population. Distance is calculated in the trait space

of the population over the generations varies. The diversity in the L-system population converges to zero quite quickly in all traits (bottom row of Fig. 9), except the one we select for: speed (Fig. 6d). In Fig. 10 we can confirm an overall loss of phenotypic diversity for L-system experiments. A visual inspection of the final populations of both setups (two sample runs are shown in Fig. 11) confirms that phenotypic diversity dropped considerably in the L-system experiments, while it's still observable in our tree-based representation. An overall decrease in phenotypic diversity in artificial evolutionary systems is often observed when evolution is converging to a solution and it is caused by a corresponding loss of genotypic diversity. A change in genotypic diversity can also explain the change in heritability we measured. A similar pattern can be observed in the tree-based representation, but the changes in heritability and diversity are visible on a smaller scale, for fewer generations and smaller changes in values.

6 Discussion

Interestingly, we observed an unexpected overall increase in heritability for the L-system, as the value for heritability started relatively low and increased over generations. We hypothesize that selection is responsible for this effect. The mating selection is probably slowly excluding all individuals that present highly unpredictable gene sequences. These gene sequences can cause very poor offspring to be generated from very fit parents and vice-versa. Slowly the selection process would pick up the unpredictable gene sequences in their low-fitness state and select them out of the next generation. This effect throughout many generations would explain a decrease in diversity and an increase in narrow-sense heritability, as only the predictable gene sequences consistently survive across multiple generations and predictable gene sequences cause high narrow-sense heritability in the population by definition.

Another pattern that we can observe is that heritability becomes highly unstable in later stages of evolution for L-system experiments; this is especially obvious in Fig. 9a. The explanation for this effect can be found by looking at the diversity (Fig. 9d). A decrease in diversity for a trait is an indicator of a decrease in overall genotypic diversity, which implies that the evolutionary process is not exploring the search space any more and all solutions are very similar. But to compute an accurate estimation of heritability we need high diversity of the population, otherwise we are computing the linear regression of a concentrated cloud of points, as shown in Fig. 12. By contrast, we deduce that the tree-based experiments are exploring the evolutionary space and having difficulties in exploiting the solution space. With this knowledge we can estimate that the tree-based representation needs a lower mutation rate and a stronger selection pressure to be able to exploit the solution space. This was to be expected because we choose very high mutation rates and a relaxed selection mechanism. On the contrary L-system seems to drop diversity after generation 20, which is pretty early. This is especially noticeable for trait which we don't select for (Fig. 9d–f). Our L-system representation seems to have difficulties to explore even from the early phases of evolution, despite using evolutionary parameters that encourage exploration. This suggests that an evolutionary process using our L-system representation has a hard time escaping a few local optima found at the beginning, and probably needs some additional elements to encourage exploration.

A final consideration can be done on the values that we measured. Real world biological experiments measure heritability values ranging from 0.30 [43] to 0.46 [44]. Instead the values we measured (Table 2) are extremely high in comparison. This difference can be explained by either one of two causes: it could be because our genotype is still quite simple when compared to real life genotypes or it could be because we have a very simple and stable environment. We reference a study on twins [49] which lived their life in very similar environments. The calculated estimates for narrow-sense heritability (h^2) in this study are very high, some as high as 0.712, comparable to some of our own results. This observation seems to

Fig. 11 Snapshot of all individuals in the last generation of two sample runs, one for each genetic representation. Images are 2D projections of the robots
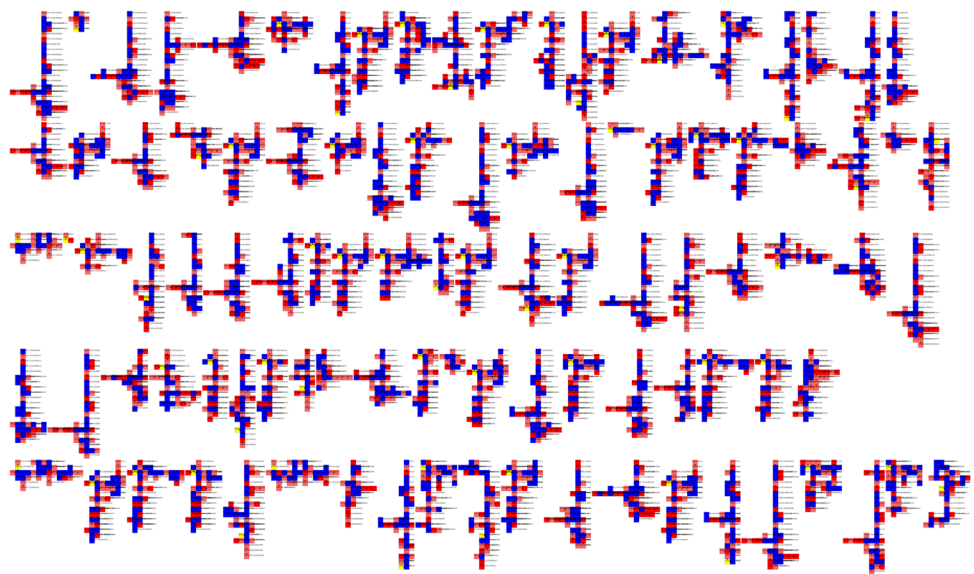
(a) Direct–tree representation, last generation
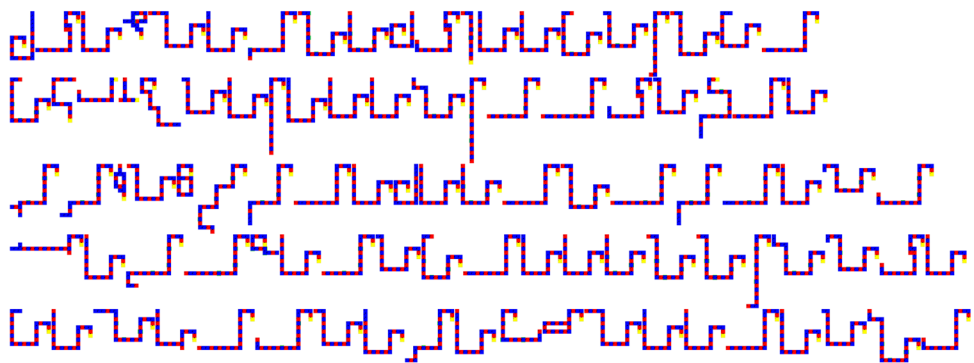
(b) L-system representation, last generation

validate our idea that the simple environment also played a role in our unexpectedly high values of heritability.

7 Concluding remarks

In this paper we introduced the biological notion of heritability as a novel tool to study representations in evolutionary robotics. Heritability captures the correlation between a quantifiable phenotypical trait measured in the parents and the one measured in the offspring. In our experiments we show that heritability can be a useful tool in evolutionary robotics to support the genotype design process. We used this novel tool to tackle the bootstrapping problem, because it reveals how exploratory a system is during the initial phases of evolution. We observed how changes in heritability could correlate to changes in diversity towards the course of evolution; i.e. in our tree-based system diversity and heritability seem to stabilize, while, in the L-system

experiments, diversity drops and the estimated heritability increases at first, followed by high instability. We related the different rates of heritability and diversity to more exploratory or exploitative evolutionary algorithms and how these concepts seem to be intertwined; i.e. we observed how high narrow-sense heritability corresponds to an exploratory evolutionary system, and how low heritability is observed in more an exploitative evolutionary system, caused by a greater epistasis effect in the representation. Importantly, this analysis can be performed only within the first few generations (in our case, 50), during the transitory phase of the evolutionary process.

Heritability proved to be a helpful tool to evaluate the shape and smoothness of the search-space, considering the landscape of both fitness and other phenotypic traits. Tree-based experiments converge to solutions where robots still retain significant morphological diversity, meaning the local optima found by evolution in the search-space is a smooth wide hill. This was expected by a genetic
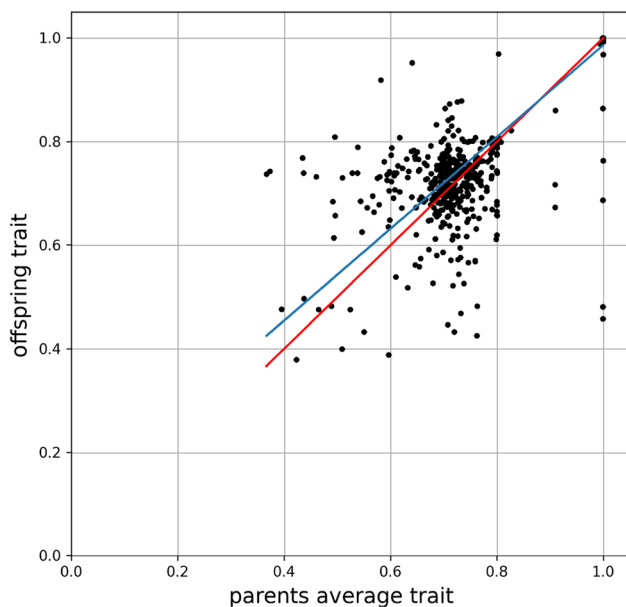
Fig. 12 Scatter plot and heritability for balance of the last generation for the L-system configuration. The x axis is the average balance of the parents and the y axis is the value of balance for the offspring. The blue line represents the linear regression of these values. A steeper line indicates a higher level of heritability, to a theoretical maximum of 1 (45° slope). The red line is a 45° reference line. The linear regression results in almost a 45° line because there are many overlapping points in the plot that are indistinguishable but influence the regression results (colour figure online)

representation that is mostly made of additive genes. On the contrary, L-system experiments converged to a single morphological solution with little-to-none diversity, and by observing the change in heritability and diversity we can determine that solutions are unlikely to explore other peaks, probably because they are either too distant, too narrow or not good enough compared to the solution found. This is an indication that the L-system is a genetic representation that contains lots of epistatic effects, meaning many genes need to align to be able to create a positive effect on the phenotype.

In possess of this knowledge, we will transition to a tree-based direct representation in future work where the focus will be on evolution of morphological traits interacting with other elements, because differences in phenotypical traits will be more tangible. In other setups, where the interest is in studying complicated gene interaction, including epistatic effects, L-system is a good candidate.

This paper also highlights how the notion of heritability draws attention to major discrepancies between biological and artificial evolutionary systems. Crucially, in biology computing a value of heritability requires measuring qualities of the traits at the phenotype level. However, phenotypic traits in biology are a result of several processes and factors:

recombination, mutation, embryonic development, early life development, behavioural and morphological lifetime adaptation. Many of these processes can be influenced by complex environments. In addition, biological systems have and retain high genotypic variation.

In contrast, artificial systems are extremely simplified. In particular, in our evolutionary system, individuals develop before they can have any interaction with the environment. When individuals are placed in the environment, interactions are extremely simple and consistent across all individuals. Additionally, individuals have no morphological or behavioural adaptation systems at their disposal during their lifetime. Artificial systems are also characterized by a generally lower genotypic diversity and a tendency to decrease even more over the course of artificial evolution, when the population converges to a useful solution.

For the future of evolutionary robotics, adding the above-mentioned biological elements would be very interesting, and some efforts have already been done in this direction. In [50, 51] we find efforts to introduce learning systems that enable individuals to adapt to their environment during their lifetime. Some work can also be found studying how to design a system where the genotype-phenotype mapping can be influenced by the environment [47]. However, these additions to the evolutionary process are non-trivial: research on these more complicated and realistic systems require a substantial increase in computational cost required from the evolutionary process. This results in slower iterations and difficulties in the design and parameter tuning processes, especially if one wanted to study these processes in combination.

The development of artificial life through artificial evolution is still in its infancy, and lot of work in the above directions and beyond could be done. Still, both for the current system complexity as well as for the one of future systems, in our view measuring heritability will be a useful tool that greatly increases our understanding of the relationship between phenotypes and genotypes.

Declarations

References

1. Cliff D, Husbands P, Harvey I (1993) Explorations in evolutionary robotics. Adapt Behav 2(1):73–110
2. Nolfi S, Floreano D (2000) Evolutionary robotics: the biology, intelligence, and technology of self-organizing machines. MIT Press, Cambridge
3. Floreano D, Husbands P, Nolfi S (2008) Evolutionary robotics. In: Siciliano B, Khatib O (eds) Springer handbook of robotics. Springer, Berlin, pp 1423–1451
4. Hupkes E, Jelisavcic M, Eiben AE (2018) Revolve: a versatile simulator for online robot evolution. In: International conference on the applications of evolutionary computation, pp 687–702
5. Griffiths A, Miller J, Dea S (2000) Quantifying heritability. An introduction to genetic analysis
6. Wray N, Visscher P (2008) Estimating trait heritability. Nat Educ 1(1):29
7. Doebeli M (1996) A quantitative genetic competition model for sympatric speciation. J Evol Biol 9(6):893–909
8. Ellers J, Slabbekoorn H (2003) Song divergence and male dispersal among bird populations: a spatially explicit model testing the role of vocal learning. Anim Behav 65(4):671–681
9. Thibert-Plante X, Hendry AP (2011) The consequences of phenotypic plasticity for ecological speciation: plasticity and ecological speciation. J Evol Biol 24(2):326–342
10. Lorch PD, Proulx S, Rowe L, Day T (2003) Condition-dependent sexual selection can accelerate adaptation. Evol Ecol Res 5(6):16
11. Petrie M, Roberts G (2007) Sexual selection and the evolution of evolvability. Heredity 98(4):198–205
12. Kokko H, Brooks R, McNamara JM, Houston AI (2002) The sexual selection continuum. Proc R Soc Lond Ser B Biol Sci 269(1498):1331–1340
13. Lande R, Arnold SJ (1983) The measurement of selection on correlated characters. Evolution 37(6):1210–1226
14. Servedio MR, Powers JM, Lande R, Price TD (2019) Evolution of sexual cooperation from sexual conflict. Proc Natl Acad Sci 116(46):23225–23231
15. Buresch T, Eiben AE, Nitschke G, Schut MC (2005) Effects of evolutionary and lifetime learning on minds and bodies in an artifical society. In: 2005 IEEE congress on evolutionary computation, vol 2, pp 1448–1454
16. Goddard ME, Kemper KE, MacLeod IM, Chamberlain AJ, Hayes BJ (2016) Genetics of complex traits: prediction of phenotype, identification of causal polymorphisms and genetic architecture. Proc R Soc B Biol Sci 283(1835):20160569
17. Sailer ZR, Harms MJ (2017) High-order epistasis shapes evolutionary trajectories. PLoS Comput Biol 13(5):1005541
18. Cho A (2014) The accidental roboticist. Science 346(6206):192–194
19. Long J (2012) Darwin's devices: what evolving robots can teach us about the history of life and the future of technology
20. Waibel M, Floreano D, Keller L (2011) A quantitative test of Hamilton's rule for the evolution of altruism. PLoS Biol 9(5):1000615
21. Floreano D, Keller L (2010) Evolution of adaptive behaviour in robots by means of Darwinian selection. PLoS Biol 8(1):1000292
22. Floreano D, Mitri S, Magnenat S, Keller L (2007) Evolutionary conditions for the emergence of communication in robots. Curr Biol 17(6):514–519
23. Gajewski A, Clune J, Stanley KO, Lehman J (2019) Evolvabilityes: scalable and direct optimization of evolvability. In: Proceedings of the genetic and evolutionary computation conference, pp 107–115
24. Lehman J, Stanley KO (2011) Improving evolvability through novelty search and self-adaptation. In: 2011 IEEE Congress of Evolutionary Computation (CEC). IEEE, pp 2693–2700
25. Mengistu H, Lehman J, Clune J (2016) Evolvability search: directly selecting for evolvability in order to study and produce it. In: Proceedings of the genetic and evolutionary computation conference 2016 (GECCO'16). Association for Computing Machinery, New York, pp 141–148
26. Auerbach J, Aydin D, Maesani A, Kornatowski P, Cieslewski T, Heitz G, Fernando P, Loshchilov I, Daler L, Floreano D (2014) RoboGen: robot generation through artificial evolution. In: Artificial Life 14: proceedings of the fourteenth international conference on the synthesis and simulation of living systems, pp 136–137
27. Jelisavcic M, De Carlo M, Hupkes E, Eustratiadis P, Orlowski J, Haasdijk E, Auerbach JE, Eiben AE (2017) Real-world evolution of robot morphologies: a proof of concept. Artif Life 23(2):206–235
28. De Carlo M, Ferrante E, Eiben AE (2020) Comparing indirect encodings by evolutionary attractor analysis in the trait space of modular robots. In: Proceedings of the 2020 genetic and evolutionary computation conference (GECCO 2020), pp 73–74
29. Hooper SL (2000) Central pattern generators. Curr Biol 10(5):176–179
30. Marder E, Bucher D (2001) Central pattern generators and the control of rhythmic movements. Curr Biol 11(23):986–996
31. Ijspeert AJ (2008) Central pattern generators for locomotion control in animals and robots: a review. Neural Netw 21(4):642–653
32. Bucher D, Haspel G, Golowasch J, Nadim F (2015) Central pattern generators. Wiley, Hoboken, pp 1–12
33. Lan G, Jelisavcic M, Roijers DM, Haasdijk E, Eiben AE (2018) Directed locomotion for modular robots with evolvable morphologies. In: Parallel problem solving from nature—PPSN XV, pp 476–487
34. Lan G, De Carlo M, van Diggelen F, Tomczak JM, Roijers DM, Eiben AE (2021) Learning directed locomotion in modular robots with evolvable morphologies. Appl Soft Comput 111:107688
35. Carlo MD, Zeeuwe D, Ferrante E, Meynen G, Ellers J, Eiben AE (2020) Robotic task affects the resulting morphology and behaviour in evolutionary robotics. In: 2020 IEEE symposium series on computational intelligence (SSCI), pp 2125–2131
36. Jelisavcic M, Glette K, Haasdijk E, Eiben AE (2019) Lamarckian evolution of simulated modular robots. Front Robot AI 6:1–15
37. Lindenmayer A (1968) Mathematical models for cellular interactions in development II. Simple and branching filaments with two-sided inputs. J Theor Biol 18(3):300–315
38. Stanley KO, D'Ambrosio DB, Gauci J (2009) A hypercube-based encoding for evolving large-scale neural networks. Artif Life 15(2):185–212
39. Gauci J, Stanley KO (2010) Autonomous evolution of topographic regularities in artificial neural networks. Neural Comput 22(7):1860–1898
40. Miras K, Eiben AE (2019) The impact of environmental history on evolved robot properties. In: ALIFE 2019, pp 396–403
41. Stanley KO (2007) Compositional pattern producing networks: a novel abstraction of development. Genetic Program Evol Mach 8(2):131–162
42. Miras K, Eiben AE (2019) Effects of environmental conditions on evolved robot morphologies and behavior. In: Proceedings of the

2019 genetic and evolutionary computation conference (GECCO 2019), pp 125–132

43. Dochtermann NA, Schwab T, Anderson Berdal M, Dalos J, Royauté R (2019) The heritability of behavior: a meta-analysis. J Hered 110(4):403–410

44. Mousseau TA, Roff DA (1987) Natural selection and the heritability of fitness components. Heredity 59(2):181–197

45. Lynch M, Walsh B et al (1998) Genetics and analysis of quantitative traits, vol 1. Sinauer, Sunderland

46. Auerbach JE, Bongard JC (2014) Environmental influence on the evolution of morphological complexity in machines. PLoS Comput Biol 10(1):1–17

47. Miras K, Ferrante E, Eiben AE (2020) Environmental regulation using plasticoding for the evolution of robots. Front Robot AI 7:107

48. De Carlo M, Zeeuwe D, Ferrante E, Meynen G, Ellers J, Eiben AE (2020) Influences of artificial speciation on morphological robot evolution. In: 2020 IEEE symposium series on computational intelligence (SSCI). IEEE, Canberra, pp 2272–2279

49. Polderman TJC, Benyamin B, de Leeuw CA, Sullivan PF, van Bochoven A, Visscher PM, Posthuma D (2015) Meta-analysis of the heritability of human traits based on fifty years of twin studies. Nat Genet 47(7):702–709

50. Miras K, De Carlo M, Akhatou S, Eiben AE (2020) Evolving-controllers versus learning-controllers for morphologically evolvable robots. In: Castillo PA, Jiménez Laredo JL, Fernández de Vega F (eds) Applications of evolutionary computation, vol 12104, pp 86–99

51. Gupta A, Savarese S, Ganguli S, Fei-Fei L (2021) Embodied intelligence via learning and evolution. Nat Commun 12(1):5721