



Evolving and generalising morphologies for locomoting micro-scale robotic agents

Matthew Uppington^{1,2} · Pierangelo Gobbo³ · Sabine Hauert^{1,4} · Helmut Hauser^{1,4}

Received: 1 November 2022 / Revised: 11 April 2023 / Accepted: 12 April 2023 / Published online: 1 July 2023
© The Author(s) 2023

Abstract

Designing locomotive mechanisms for micro-scale robotic systems could enable new approaches to tackling problems such as transporting cargos, or self-assembling into pre-programmed architectures. Morphological factors often play a crucial role in determining the behaviour of micro-systems, yet understanding how to design these aspects optimally is a challenge. This paper explores how the morphology of a multi-cellular micro-robotic agent can be optimised for reliable locomotion using artificial evolution in a stochastic environment. We begin by establishing the theoretical mechanisms that would allow for collective locomotion to emerge from contractile actuations in multiple connected cells. These principles are used to develop a Cellular Potts model, in order to explore the locomotive performance of morphologies in simulation. Evolved morphologies yield significantly better performance in terms of the reliability of the travel direction and the distance covered, compared to random morphologies. Finally, we demonstrate that patterns in evolved morphologies are robust to small imperfections and generalise well to larger morphologies.

Keywords Morphology · Artificial evolution · Micro-scale · Locomotion

1 Introduction

Locomotion is a crucial ability for micro-scale systems, particularly in applications that involve the transportation of cargos, or self-assembling in pre-programmed architectures.

Sabine Hauert and Helmut Hauser contributed equally to this work.

✉ Matthew Uppington
mu15531@bristol.ac.uk

Pierangelo Gobbo
PIERANGELO.GOBBO@units.it

Sabine Hauert
sabine.hauert@bristol.ac.uk

Helmut Hauser
helmut.hauser@bristol.ac.uk

¹ Bristol Robotics Laboratory, Farscope CDT, University of West England, 100190 Bristol, UK

² Department of Aerospace Engineering, University of Bristol, 10587 Bristol, UK

³ Department of Chemical and Pharmaceutical Sciences, University of Trieste, 610101 Trieste, Italy

⁴ Department of Engineering Mathematics, University of Bristol, Bristol, UK

The ability to produce micro-scale systems composed of multi-cellular units [1, 2], and functionalise them [3, 4], have seen several advancements in recent years. Breakthroughs in the development and understanding of contractile proto-cells and light-based testing environments bring the promise of light-controllable, mobile micro-robots closer to reality, [5–8]. Yet, understanding how to design locomotion from multi-cellular morphologies is still an open question. Morphology in this context refers to the specific form or structure of a robot or group of robots [9, 10]. Artificial evolution has been used to search for optimal robot morphologies in many studies [11, 12].

Exploring the potential abilities of reactive, micro-scale systems in simulation is a vital first step towards eventually deploying new micro-robots in real-world applications. In this paper, we consider a system of micro-scale cellular units that can expand and contract in response to a light-based stimulus, as shown in Fig. 1. This is inspired primarily by the works of Gobbo et al. [5] and Downs et al. [4], where similar forms of actuation are presented with proto-cells and hydrogel structures respectively. Using light as a control medium is advantageous for independently and precisely targeting micro-scale agents [13]. This can be achieved with open-source technology such as the Dynamic Optical

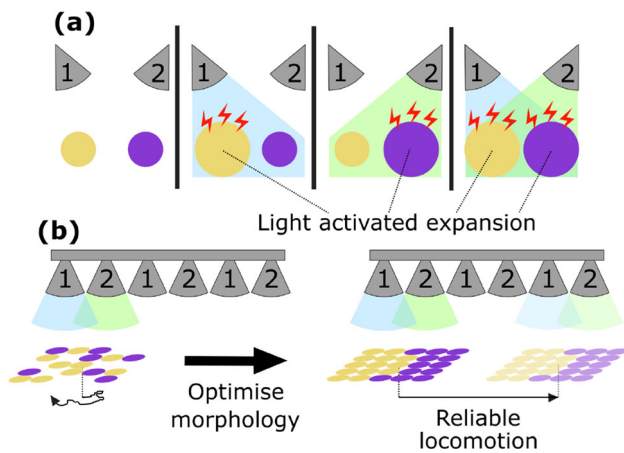


Fig. 1 (a) Conceptual illustration of micro-scale, light-reactive, contractile cellular units – different cells (yellow / purple) expand when exposed to different wavelengths of light. (b) By using artificial evolution in a stochastic simulator, we seek to optimise the morphology of a connected group of cells to allow locomotion in a reliable direction in a stochastic environment

Micro Environment (DOME) [7], which offers closed-loop control of projected light patterns for illuminating and imaging micro-scale systems.

The ability to control multiple agents reliably and independently in a given environment is well studied at the macro scale [14, 15]. Having access to multiple, independently controllable robots can help to parallelise tasks, or open up new types of cooperative and higher order behaviours. This can be of great benefit for application scenarios such as simultaneous capture of multiple target cargos [16]. However, transferring the vast array of abilities that multi-robot platforms are capable of to real systems at the micro-scale is a challenge. As advancements are made in developing new types of reactive micro materials [5], it is of interest to investigate their potential for instigating multi-robotic behaviours.

As a first step, mobile morphologies for a collection of contractile cells were learned in simulation using artificial evolution, as presented by Uppington et al. [17] (© 2022 IEEE. Reprinted, with permission, from M. Uppington, P. Gobbo, S. Hauert and H. Hauser). Here, we expand further on this work by also showing how patterns in our learned morphologies can be generalised to create new morphologies with desired directions of locomotion and increased reliability (see Section 4.3). We also investigate how the size of these morphologies, in terms of the number of constituent protocells, impacts the locomotive performance (see Section 4.3), and we elaborate further on the mathematical premise of the actuation model (see Section 3.1).

In the next section we summarise some previous works related to how locomotion has been achieved with micro-robots, how their morphologies can be optimised and the available methods for simulating them. We then begin Sec-

tion 3 by considering the minimal requirements for achieving locomotion in light-activated contractile cells. These principles are translated into a Cellular Potts model using CompuCell3D as a simulator for micro-scale cellular units in 2 dimensions. By using artificial evolution, we search for optimised morphologies that display reliable locomotion. The performance metrics and hyper-parameters used for our artificial evolution trials are outlined in Sections 3.2 and 3.3 respectively. In Section 4.1 we first show the behaviour of the simple, two-cell morphology we describe in Section 3.1, for comparison with the behaviours of evolved morphologies, which are presented in Section 4.2. We investigate how patterns in the evolved morphologies can be generalised to design new morphologies with predictable directions of locomotion, and explore how well these morphologies scale to larger numbers of protocells in Section 4.3. We conclude with a discussion of our results and an outline of potential avenues for future research.

2 Related work

Designing locomotive micro-scale robots with optimal morphologies has been tackled in many ways and in a variety of contexts. We begin in Section 2.1 by broadly exploring how micro-scale locomotion has been implemented in previous studies. We discuss the unique advantages of light-based control, as well as some limitations of existing systems that use this method. Then, Section 2.2 explores the application of artificial evolution to optimising morphologies. Finally, we consider some potential simulation platforms with which to design and implement our virtual experiments in Section 2.3.

2.1 Micro-robot locomotion

Locomotion has been achieved in micro-robots through a variety of different methods. For example, Downs et al. presented a multi-responsive hydrogel structure that exemplifies several stimuli being used together, including light, temperature and magnetism [4]. In general, magnetic and light-based control are used, either separately or jointly, by the majority of micro-robotic systems [18]. Magnetism has been used to design many micro-scale systems whose movements can be controlled with a high degree of precision [19–21]. However, a drawback of using magnetism for locomotion is that a single control signal is applied globally via a magnetic field, which limits the extent to which multiple agents can be controlled independently. Some studies such as Xie et al. have demonstrated that magnetism can be used to operate a swarm of many micro-robots in tandem, but not with locally independent behaviours [22].

In contrast, light, when used as a control medium, can facilitate many independent signals being directed towards

many different agents. Miskin et al. demonstrated light-based control with micro-swarming robots that locomote when top-mounted photo-voltaic receptors are targeted with light lasers [13]. As a result, due to the high resolution that is achievable with laser targeting, each individual micro-robot in the swarm can be controlled independently. However as lasers only actuate either the front or back pairs of legs on each micro-robot, locomotion is somewhat limited to moving forwards or backwards. Lv et al. use light to manipulate individual micro-particles, or aggregate large groups of micro-particles into a local area with Light Activated Marangoni Tweezers (LAMT), although only in the presence of a photosurfactant solution [23]. Zeng et al. also use light to control a micro-walker, though the locomotive behaviours are also partly coupled with the surface type of the environment [24].

2.2 Optimisation of morphologies

With an ever increasing arsenal of materials to functionalise micro-scale systems, the task of optimising their design for a particular function, such as locomotion, is challenging. Even with just a small set of building blocks, the range of possible combinations to construct morphologies makes it difficult to search efficiently for optimal solutions. One solution is artificial evolution, which is a machine learning technique that has a long history of being applied to optimising the design and behaviour of robots [25]. Mautner and Belew showed how robot morphology could be evolved in parallel with a controller in simulation [11]. Since then, a similar concept of simultaneously evolving morphology and controller has been applied to physical robots [12].

One challenge is being able to bridge the reality gap between morphologies evolved in simulation and morphologies produced in the real world. In the microscopic domain, Kriegman et al. presented a solution for a scalable process for designing and producing mobile morphologies using artificial evolution [2]. Morphologies were optimised for locomotion by measuring performance in terms of total displacement achieved in simulation during a given time. Whilst this does produce mobile micro-robots, called Xenobots, their trajectories cannot be controlled by external factors (such as light). Kriegman et al. also note the importance of testing high performing (simulation) morphologies in the presence of noise, in order to filter out which simulated morphologies are likely to perform well in the real world [2].

2.3 Cellular simulators

The choice of simulation environment has a significant impact on how systems can be represented and, therefore, how they can be optimised with artificial evolution. The BSim simulator, described by Gorochoowski et al., allows collective behaviours of agent-based, micro-scale systems

to be modelled in complex environments, though it is specifically designed to model bacterial cells [26]. Kriegman et al. defined morphologies for Xenobots in simulation using a voxel grid [2]. Voxel grids allow virtual organisms to be constructed out of 3-dimensional arrays of customisable, cubic building blocks [27]. While voxel grids are an intuitive and effective strategy, other simulation environments are better tailored to modelling micro-scale systems.

Cickovski et al. introduced *CompuCell3D* as a flexible simulation environment for micro-scale, cellular systems [28]. The underlying model is Cellular Potts; a powerful, generalised framework for simulating micro-scale systems using a pixel grid [29, 30]. Cellular Potts models are intrinsically stochastic meaning that the need to introduce noise, as highlighted by Kriegman et al., is satisfied naturally. In addition, *CompuCell3D* provides a powerful Python based interface, with built-in optimisation tools to manage CPU usage and modular run-time functionalities.

These are desirable features for balancing computational speed versus complexity when running demanding artificial evolution algorithms and, therefore, was chosen as our simulation environment. The next section describes how our protocell-like contractile cells were simulated using *CompuCell3D*.

3 Methods

3.1 Contractile cell model

We use *CompuCell3D* to define a Cellular Potts model of a connected group of generalised cells. In Cellular Potts models, cells are represented by a region of pixels within a grid and the simulation is progressed incrementally in Monte Carlo Steps (MCS). In each MCS, each pixel in the grid attempts to copy their value to a nearby pixel with some probability of success. The core principle behind Cellular Potts models is that interactions between entities can be defined in terms of energy equations that determine the probability of successful pixel copies [28–30]. These energy equations most commonly take the form of a quadratic equation,

$$E_i(v_i) = M_i(v_i - C_i)^2, \quad (1)$$

where v_i is some variable that is measured in the system (such as the volume of a given cell) and $E_i(v_i)$ is the energy term related to this variable. The parameter C_i represents the equilibrium value that the variable v_i is driven towards, and the coefficient M_i determines how quickly energy is accumulated as v_i deviates from its equilibrium value. In general, pixel copies that act to reduce the total energy in the system, $\sum_i E_i(v_i)$, have a higher probability of success. By adjusting

the values of M_i and C_i in Eq. 1 for certain variables over time, dynamic and reactive behaviours can be simulated.

Our model is inspired by protocells introduced by Gobbo et al. [5], that have since been developed further into protocellular materials suspended in water (see Galanti et al. [6]). In this paper, we use “cell” to refer to a contractile protocellular agent, rather than a cell in a biological context. We assume that each cell is able to increase its radius up to its maximum when exposed to a stimulating control signal. So, we introduce an energy equation to control the length of a cell in its primary (longest) axis and secondary (orthogonal to primary) axis. We take a hypothetical $10\mu\text{m}$ wide protocell at room temperature as inspiration for the width of a cell when unstimulated, as introduced by Gobbo et al. [5]. The stochasticity in Cellular Potts models will mean that the exact shape of a given cell will change slightly with each simulation step. To ensure a cell does not deviate unrealistically from roughly uniform shapes (circular in 2D), additional energy equations are included to govern the perimeter and surface area of the cell as a function of its radius. To ensure neighbouring cells remain connected, the ideal distance between their centres of mass is defined as the sum of their radii (see Eq. 4 further down). Note that these connections, referred to as focal point plasticity links, are loose, meaning that cells can freely rotate about other cells, unless otherwise constrained. In addition, to prevent cells from being absorbed into each other, energy equations are introduced to apply a small repulsive force between pixels belonging to different cells. The balance of these interactions allows cells to stay physically connected at their surface, whilst each maintaining integrity and roughly uniform shapes.

To understand how a system consisting of these types of cellular units could locomote, we start by considering the simplest possible case, i.e. a connected group of just two cells. If each cell were assumed to be perfectly spherical and surrounded by a viscous, stationary fluid, then their dynamics can be modelled using Stoke’s law [31],

$$F(t) = 6\pi\mu r(t) \frac{dx(t)}{dt}, \quad (2)$$

where F is the force experienced by the cell, μ is the dynamic viscosity of the fluid, $r(t)$ is the radius of the cell and $x(t)$ is the position of the cell relative to the fluid. For two cells, at positions $x_1(t)$ and $x_2(t)$, with corresponding radii $r_1(t)$ and $r_2(t)$, their one-dimensional motion can be described using Eq. 2 as,

$$6\pi\mu r_1(t) \frac{dx_1(t)}{dt} + 6\pi\mu r_2(t) \frac{dx_2(t)}{dt} = 0. \quad (3)$$

Here, it is assumed that the cells are not influenced by any external forces, so the resultant is 0. Due to symmetry, it is assumed that motion occurs only along the one dimensional

axis passing through the centres of both cells. For two agents that are connected to each other, an additional constraint can be applied to ensure their surfaces stay in contact,

$$|x_2(t) - x_1(t)| = r_1(t) + r_2(t). \quad (4)$$

Note that cells have no direct control over their position; a cell responds to a control signal by varying its radius only (see Fig. 1). So, isolated agents that undergo isotopic expansions and contractions would not be expected to locomote. In order to move, cells will need to exploit the local interactions that arise from their physical connection as they expand and contract. The trajectory traced by the morphology’s geometric centre, $x_C(t) = \frac{1}{G} \sum_{i=1}^G x_i(t)$, is indicative of its locomotive potential. Here, $G = 2$ for our example with two cells. Using this definition, the velocity of the morphology can be found by solving Eqs. 3 and 4. Hence the velocity of our two-cell morphology is

$$\frac{dx_C(t)}{dt} = \frac{1}{2} \frac{r_1(t) - r_2(t)}{r_1(t) + r_2(t)} \left(\frac{dr_1(t)}{dt} + \frac{dr_2(t)}{dt} \right). \quad (5)$$

It is assumed here, without loss of generality, that cell 2 is positioned to the right of cell 1, and that velocities are measured as being positive to the right. Eq. 5 shows that two connected cells cannot locomote if their radii remain equal, even when expanding or contracting, $\forall t : r_1(t) = r_2(t) \Rightarrow \frac{dx_C(t)}{dt} = 0$. Eq. 5 also shows no motion can be expected if one cell is expanding at the same rate as the other cell is contracting, $\forall t : \frac{dr_1(t)}{dt} = -\frac{dr_2(t)}{dt} \Rightarrow \frac{dx_C(t)}{dt} = 0$. We assume that expansions and contractions are reversible and occur at the same rate. Therefore, in order to locomote, it is necessary that cells do not maintain equal radii and that expansions and contractions do not always occur simultaneously.

Figure 2 shows an example of a simple pair of asynchronous control signals that minimise the time that cells spend with equal radius whilst also avoiding simultaneous expansions and contractions. In a practical implementation with protocells, these control signals could be transmitted with different wavelengths of light (compare Fig. 1). Alternatively, different cells could be individually stimulated with the same wavelength of light at different times, which is possible, for example, with the DOME technology [7].

Here, a morphology that is optimised for locomotion should exploit the actuation shown in Fig. 2 to be able to travel the greatest possible distance. However, to maximise transferability to the real-world, morphologies should also be able to overcome the effects of noise - their locomotion should be repeatable. To evaluate the performance of morphologies, we define a fitness function for reliable locomotion in the next section.

3.2 Performance metric

The goal of optimisation is to find morphologies that can locomote the greatest distance most reliably. Therefore, the locomotive performance of a morphology can be quantified using the displacement of the morphology’s geometric centre after a certain number of simulation steps. To drive locomotion, the radii of each cell in a given morphology is varied according to their type, as shown in Fig. 2 for 400 MCS; this will correspond to one cycle of actuation. By simulating a morphology for T actuation cycles, we can obtain a vector describing the final position of the morphology’s geometric centre, with respect to its initial position, \mathbf{p} . The direction of the morphology’s motion is indicated by the angle that this final position vector makes with the horizontal axis, θ . Given the stochastic nature of Cellular Potts models, tracking the average performance not just over multiple actuation cycles, but also over multiple independent simulations provides a more valuable description of the morphology’s expected behaviour. So, running K simulations of a given morphology provides us with a set of independently sampled displacement vectors, $\mathbf{p}_1, \dots, \mathbf{p}_K$ with corresponding directions $\theta_1, \dots, \theta_K$. To optimise only over the total distance covered we could simply take the average magnitude of these vectors, i.e., $D = \frac{1}{K} \sum_{j=1}^K |\mathbf{p}_j|$. However, here in order to encourage morphologies also to travel in a reliable direction, we want to minimise the variation over the direction (i.e., the spread / deviation in θ). Therefore, we define the following

fitness function,

$$F = \frac{1}{T} \left| \sum_{j=1}^K \frac{\mathbf{p}_j}{K} \right|. \tag{6}$$

We divide by T to average over the number of actuation cycles that morphologies are simulated for. By taking the magnitude of the averaged vector (instead of the average of the vector magnitudes), morphologies are penalised for travelling in random directions, but rewarded for travelling further in similar directions in each run. Note that this fitness function is bounded by the average distance covered, $F \leq \frac{D}{T}$, since $\left| \sum_{j=1}^K \mathbf{p}_j \right| \leq \sum_{j=1}^K |\mathbf{p}_j|$ via the triangle inequality. This maximum limit is approached as the standard deviation in travel directions $\sigma = \sqrt{\frac{1}{K} \sum_{j=1}^K (\theta_j - \bar{\theta})^2}$ approaches zero, where $\bar{\theta} = \frac{1}{K} \sum_{j=1}^K \theta_j$ is the average direction. Thus, higher performing morphologies are identified as those that travel larger distances on average and have smaller deviations in travel directions. With a fitness function defined, we outline the setup for our evolutionary algorithm in the next section.

3.3 Evolutionary setup

To explore a large space of possible morphologies, analytical methods quickly become infeasible. For a given arrangement of G cells that can be in one of S states, the number of possible morphologies scales with S^G . Whilst intuition might suggest morphologies that could perform well, having a systematic approach provides a tool to explore a wider range of morphologies. However, simply using exhaustive search to find optimal morphologies within such a vast set of possible solutions is not feasible either, due to being prohibitively time-consuming. Artificial evolution provides a solution, by starting with a batch of random morphologies, called the initial population or first generation. These morphologies are evaluated using the performance metric defined in Section 3.2 (see Eq. 6). This produces a single numerical value for each morphology, allowing them to be ranked from best (highest value) to worst (lowest value). A new generation can then be created by replacing poorly performing morphologies with random combinations of the better performing morphologies (crossover), sometimes with random changes (mutations). The best morphologies from one generation are carried over to the next generation unaltered. By iteratively repeating this process for subsequent generations the performance can be gradually increased over multiple generations.

To establish the search space of morphologies, we arrange the initial positions of cells into a 5×5 Cartesian formation, as shown in Fig. 3. Each position either is occupied by a cell that expands when stimulated by one of the two active control




Fig. 2 Graph showing asynchronous control signals with a period of 400 simulation steps (MCS). The early (yellow, lower) and late (purple, middle) signals cause cells radii to oscillate between a normalised minimum value, and a maximum value. Passive cells (turquoise, upper) are not stimulated. A cell’s maximum radius is based on the scaling coefficient that would double the volume of a spherical cell. Thus $\frac{\max(r_i(t))}{\min(r_i(t))} = 2^{\frac{1}{3}} \approx 1.26$ (not shown to scale)




Fig. 3 Diagram showing how a random morphology with a 5×5 Cartesian grid structure can be represented as a genome of length $G = 25$. First, each row of cells is listed in order, starting from the top, moving down the grid (a). Each cell state in the ordered list is then represented with an index (b): ‘E’ = early actuating cell (yellow), ‘L’ = late actuating cell (purple), ‘P’ = passive cell (turquoise), and ‘N’ = no cell (white)

signals described in Fig. 2, is occupied by a passive cell that does not expand, or does not contain a cell. By considering ‘no cell’ as a fourth cell state, the size of the search space is defined by $S = 4$ and $G = 25$. This leads to $S^G = 4^{25}$ possible morphologies. Figure 3 explains how morphologies are encoded by listing rows of cell states from top to bottom. This allows the morphologies in the 5×5 Cartesian format to be uniquely expressed as genomes of length $G = 25$.

To generate an initial population of morphologies for the evolutionary algorithm, 32 randomly generated genomes were created. By experimenting with different values for K and T , we found that simulating for $T = 10$ actuation cycles, repeated $K = 4$ times, offered an appropriate balance of reliability in the fitness score verses simulation run time during evolution. The performance of a genome was calculated as a function of the sampled vectors collected from these repeated simulations according to Eq. 6. When progressing from one generation to the next, we selected the top 6 best performing genomes to survive. The worst 13 genomes were overwritten with new randomised sequences of cell states. The remaining 13 genomes were replaced with combinations of parent genome pairs (selected from the best 6) using uniform crossover. To analyse the performance of specific morphologies in more detail after evolution, we simulated morphologies for $T = 50$ actuation cycles, over $K = 10$ independent runs.

In the next section, we evaluate the performance of our simple two-cell morphology and the morphologies obtained from artificial evolution trials. to 6 physically neighbouring agents in a hexagonal formation, compared to just 4 with a cartesian. Due to the stochasticity of the simulation, the same cartesian morphology may settle in to different hexagonal formations on subsequent repeats. Hence, in order to increase

the stability of the agents initial states, the cartesian grid is rearranged to match the hexagonal lattice seen in the later stages of simulation.

4 Results

All experimental results were collected by running our simulator on an Acer Aspire E15 laptop with a 4-core, Intel I5-7200U CPU and 8GB DDR4 memory for periods of up to 6 hours; implemented using CompuCell3D v4.2.1-v4.2.4 and Python 3.

4.1 Two-cell morphology

To establish a benchmark for locomotive performance, we first simulate the two-cell morphology described in section Section 3.1. Figure 4 shows a sample of displacement vectors as well as some typical trajectories travelled by the two-cell morphology (morphology is depicted in top, left corner). While the morphology does achieve locomotion to some degree, the direction of travel appears to be highly unpredictable with a large standard deviation, $\sigma = 79.0^\circ$. The average distance travelled, $D = 6.8\mu\text{m}$, also seems relatively small, compared to the diameter of a single unstimulated cell ($10\mu\text{m}$). The chaotic nature of the morphology’s typical trajectories provides some context for these measurements, suggesting that the two-cell morphology is poorly suited to




Fig. 4 Sampled displacement vectors of the geometric centre of a two-cell morphology – one early actuating cell (yellow) and one late actuating cell (purple). Vectors were measured after $T = 50$ actuation cycles, in $K = 10$ independent runs. Three typical full trajectories of the morphology’s geometric centre are shown. The other runs are shown only by their resulting performance vectors, \mathbf{p}_j . The black scale bar shows the diameter of one unstimulated cell. As one can see, the two-cell model travels very little and with an inconsistent direction. The average distance covered was $D = 6.8\mu\text{m}$, and the standard deviation in travel direction was $\sigma = 79.0^\circ$

overcoming the inherent noise in the stochastic simulator. We therefore extend to more complex ($G = 25$) morphologies to produce more reliable locomotion in the next subsection.

4.2 Evolved morphologies

Following the evolution protocol outlined in Section 3.3 it was found that 50 generations was sufficient to allow fitness values to increase significantly. The top 20% of recorded fitness values in a sample of three typical runs of artificial evolution is shown in Fig. 5. The best performing genomes show rapid increase in fitness values for early generations with a plateau being reached by around 30 generations. At this point, evolved morphologies seem to have achieved roughly a 3 to 4 fold increase in fitness score, compared to the initial population of random morphologies. This indicates that about 50 generations is sufficient to allow the artificial evolution algorithm to converge.

Figures 6, 7 and 8 show the highest performing evolved morphologies from three independent evolution trials. In general, the trajectories of the evolved morphologies are able to travel farther than the two-cell morphology. The average distances covered by the evolved morphologies shown range between $21.7\mu\text{m} \leq D \leq 36.1\mu\text{m}$. Given that the resting height / width of the morphologies are equal to $5 \times 10\mu\text{m} = 50\mu\text{m}$, this equates to between 43.4% – 72.2% of the morphology's body length. Moreover, the measured standard deviation in travel direction is significantly lower for the evolved morphologies, ranging between $30.1^\circ \geq \sigma \geq 24.2^\circ$. This demonstrates that the evolved morphologies are able to travel more consistently in a predictable direction.




Fig. 5 Graph showing the range (vertical bars) and average of the top 20% of recorded fitness values per generation for 3 artificial evolution trials. In each generation, each genome was simulated for $T = 10$ actuation cycles, over $K = 4$ independent runs. Fitness was calculated according to Eq. 6

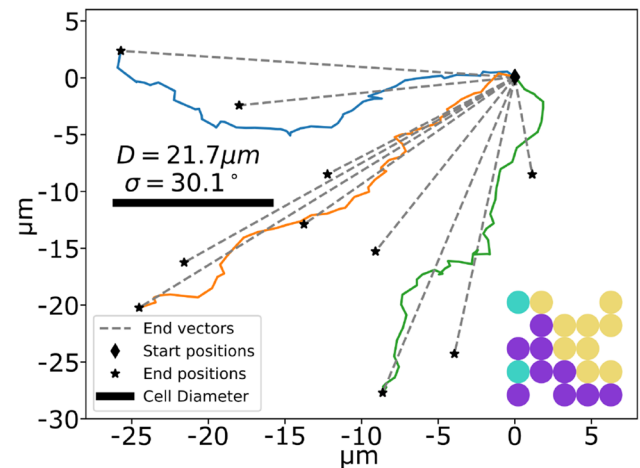


Fig. 6 Sampled displacement vectors of the geometric centre of an evolved morphology (see bottom right) from independent trial 1. Notation and simulation parameters are as described in Fig. 4

A common feature that has been seen in each of the high performing evolved morphologies is the relatively small number of ‘no cell’ and passive cell types, compared to either of the two active cell types. Figure 9 shows how the distribution of cell states in the 6 highest performing morphologies changes over the course of a typical artificial evolution trial. Again, a steady state seems to be reached after roughly 30 generations, corresponding to the formation of a plateau in fitness values (see Fig. 5) at around the same point. Similar trends were consistently found in multiple independent evolution trials, each with different, random starting configurations. This suggests that an evolutionary advantage has been found in increasing the proportion of active cells in a given morphology. Another observation is that active cells of a given type seem to cluster together with other cells of the




Fig. 7 Sampled displacement vectors of the geometric centre of an evolved morphology (see bottom left) from independent trial 2. Notation and simulation parameters are as described in Fig. 4

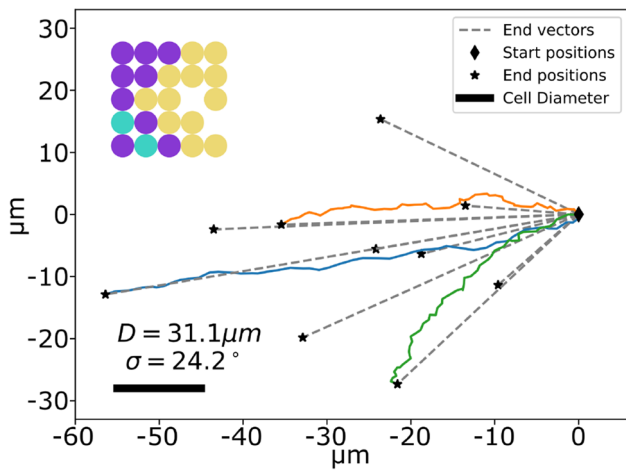


Fig. 8 Sampled displacement vectors of the geometric centre of an evolved morphology (see top left) from independent trial 3. Notation and simulation parameters are as described in Fig. 4

same type. Furthermore, the late actuating cells (purple) seem to cluster towards the side of the morphology that faces the average direction of travel. In the next subsection, we explore how these seemingly general patterns can be exploited to create new, hand-crafted morphologies with predictable locomotive behaviours.

4.3 Generalising morphologies

The morphologies learned by artificial evolution are surprisingly simple. However, the similarities between the evolved morphologies of multiple independent evolution trials aids

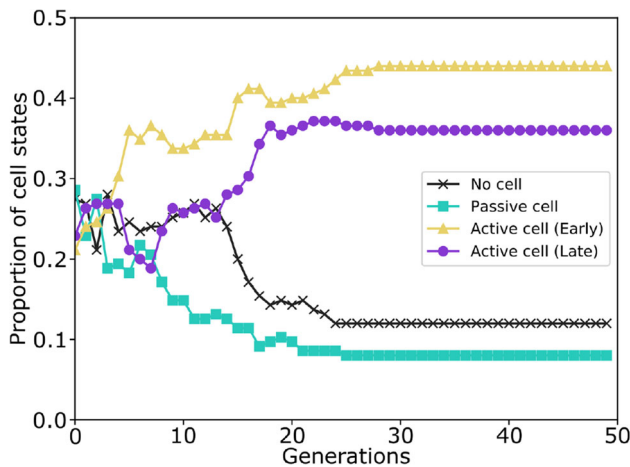


Fig. 9 Graph showing the average composition (i.e., types of cells) of the best 20% of morphologies over generations. Passive cells and “no cells” are far less common than either of the actuating cell types in later generations

further investigation through generalisation. Given the relative abundance of early and late actuating cells in the evolved morphologies, compared to passive or absent cells, we hypothesise that new morphologies consisting of only active cells could achieve similar or even better performances. If passive and ‘no cell’ states could be omitted whilst maintaining similar performance, this would greatly simplify the morphological search space. Taking the morphology depicted in Fig. 7 as a starting blueprint, we look to replace passive cell states with one of the active cell states. This process is illustrated in Fig. 10. Given that early actuating cells and late actuating cells each tend to cluster together, we determine the new cell type based on the most common active type of nearby cells. By repeating these steps for ‘no cell’ states, we arrive at a simplified morphology that consists of active cell types only.

Figure 11 compares the performance of two evolved morphologies (from Figs. 7 and 8), to the performance of two rotated variants of the hand-crafted morphology shown in Fig. 10. The alignment of late actuating cell clusters to a ‘forward’ direction in the evolved morphologies is also observed in the new hand-crafted morphologies. Furthermore the hand-crafted morphologies even display slightly improved performance compared to the original evolved morphologies. The hand-crafted morphologies travel greater distances on average, $D \geq 44.0\mu\text{m}$ (or 88.0% of a body length), and produce lower standard deviations in their travel directions, $\sigma \leq 15.7^\circ$.

These hand-crafted morphologies bare similarities to the simple two-cell morphology in terms of fundamental structure, yet their locomotion is improved in regards to both distance and reliability. The main factor that differentiates between them is the number of cells they consist of, suggesting that the stochasticity in the simulator effects larger




Fig. 10 Illustration showing how a simplified morphology (right) can be constructed from the evolved morphology from trial 2 (left). Passive cells are first replaced with the most common active state in neighbouring cells (a). Then, the same procedure is applied to ‘no cell’ states, leading to a morphology consisting of only active cells (b)

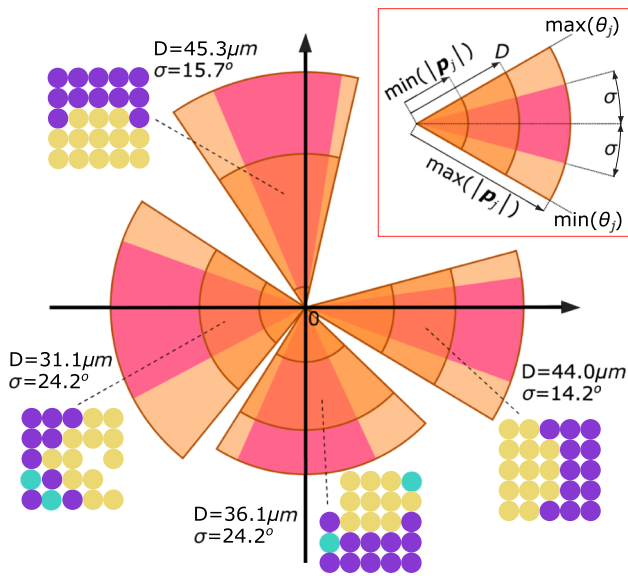


Fig. 11 Comparison between behaviours of evolved morphologies and rotated variants of the simplified morphology from Fig. 10. End vectors are sampled by simulating each morphology for $T = 50$ actuation cycles over $K = 10$ independent runs. The radii of the orange sectors, working outwards from the centre, represent the minimum, mean and maximum distances travelled by each morphology. The straight edges of the orange sectors show the total angular range of directions. The pink sector represents the standard deviation in directions, centred on the mean direction

morphologies to a lesser extent. Resilience to noise is a desirable property when seeking systems that can be transferred from simulation to real world scenarios, [2]. To further investigate how reliability scales with the number of cells, we replicate the same hand-crafted morphology with larger collections of cells. Figure 12 shows how a ‘right facing’ variant of the hand-crafted morphology (see Fig. 10) performs when doubling and tripling the morphology’s dimensions, to obtain 10×10 and 15×15 cell grids respectively. As the scale of the morphology’s body length increases, the standard deviation in travel directions indeed decreases, from $\sigma = 14.2^\circ$ for the 5×5 morphology, to $\sigma = 9.1^\circ$ for 10×10 and $\sigma = 8.3^\circ$ for 15×15 . However, the average distance travelled does not increase linearly with the morphology’s body length; travelling as far as $66.1 \mu\text{m}$ with a body length of $10 \times 10 \mu\text{m} = 100 \mu\text{m}$, but only $54.9 \mu\text{m}$ when the body length reaches $15 \times 10 \mu\text{m} = 150 \mu\text{m}$. In fact, when measured in proportion to the body lengths, the average distance travelled decreases from 88.0% of a body length for a 5×5 morphology down to 66.1% for 10×10 and 36.6% for 15×15 . This reveals that there may in fact be a trade-off between the number of body lengths that morphologies are

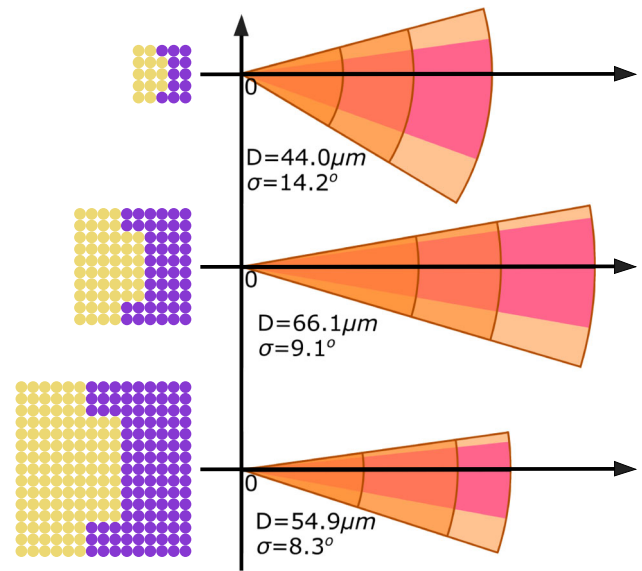


Fig. 12 Comparison between morphologies of different sizes, based on a rotated version of the simplified morphology in Fig. 10. Sector representations of sampled end vectors are plotted as in Fig. 11

able to travel (given a certain number of actuation cycles) and the reliability in their direction of travel.

5 Discussion

There are several observations that can be made about the evolved morphologies that could potentially inspire future designs for micro-scale robots. For example, we see that late actuating and early actuating cells seem to define front and back sections of evolved morphologies respectively, in terms of the direction of forward motion. This, along with the actuation cycles shown in Fig. 2, could provide a blueprint for generating locomotion in systems of connected, contractile units. In addition, the symmetries in the found morphologies suggest that the direction of travel could be changed dynamically by altering how and when each cell is stimulated. In some ways, this could be seen as analogous to inchworm-like morphologies seen in nature and other studies inspired thereby [20, 24]. Though, while these morphological features were shown to generate similar behaviours in larger cell grids, the behaviour of the smaller two-cell morphology was far less predictable. Continuations of this work could further explore the relationship between the number of cells and the locomotive performance of a given morphology, particularly for smaller grid sizes than 5×5 .

The higher performing morphologies found in this work mainly consist of active cell types, with the passive and ‘no

cell' states being far less common. This suggests that the purely structural contribution of passive or absent cells to a morphology is generally outweighed by the increased locomotive potential offered from active cell types, when it comes to generating larger and more reliable displacements. This is further supported by the comparable average distance travelled and deviation in direction displayed by hand-crafted morphologies, which consist exclusively of active cell types. The existing passive or 'no cell' states in evolved morphologies could also be interpreted as a sign of robustness to defects. For example, an active cell that loses its sensitivity to light would effectively become a passive cell, or a morphology that suffers physical damage may lose some cells by becoming detached. Future studies could look to investigate this further to examine the robustness of morphologies to different modes of failure. The omission of passive and 'no cell' states could also greatly reduce the morphological search space in subsequent artificial evolution trials, or potentially allow complexity to be introduced elsewhere, such as exploring more general actuation schemes.

Overall, we find that the evolved and our hand-crafted morphologies were able to travel significantly greater distances compared to both random morphologies (see Fig. 5), and a simple two-cell case. If the distance travelled is averaged over the number of constituent cells in the morphologies, the two-cell morphology in fact achieves the greatest average distance per cell, though at the cost of having highly chaotic trajectories. On the other hand, evolved morphologies display a much lower standard deviation in their travel direction, making their motion much more predictable. Furthermore, the distance travelled per constituent cell reduces even further when the size of morphologies are scaled up. This suggests that there may be some trade-off between locomotive potential and reliability in direction of travel when increasing the scale of the morphologies. In applications where cells, acting as or carrying a cargo, need to be delivered to a precise area, reliable travel directions would be important; whereas in applications such as area coverage, moving fewer cells greater distances in random directions may suffice, or even benefit.

The artificial evolution algorithm could be improved further in many ways. For example, rotationally symmetric morphologies could be conflated in each generation. This would reduce the search space significantly but would also incur an additional computational cost for the pair-wise comparisons between all morphologies in each generation. To investigate more complex morphological search spaces, the evolutionary algorithm could instead be supplied with larger building blocks, each consisting of multiple cells. These building blocks could themselves be inspired by the evolved

morphologies presented here or other handcrafted morphologies. Artificial evolution is a versatile tool that could also be adapted to search for alternate modes of locomotion. For instance, the fitness function could instead be redefined to search for morphologies that spin/rotate rather than translate.

We envisage future implementations of these evolved morphologies in physical systems with protocellular materials, although other photo-contractible materials could also be explored such as hydrogel structures [4]. For achieving closed loop control of these types of light reactive systems, we see the DOME as a promising and adaptable platform for future, real world experiments [8].

6 Conclusion

We have shown that artificial evolution can be used to find morphologies that allow simulated groups of connected, contractile cells to locomote consistently in the presence of noise. Hand-crafted morphologies inspired by the evolved solutions demonstrated even better performance. The patterns shown in these morphologies were also shown to generalise to different sizes of cell grids. However, whilst implementing mobile morphologies at larger scales seems to increase reliability in the direction of travel, this may come at the cost of sacrificing distance in terms of the number of body lengths that can be travelled. This work could inspire future designs for mobile, micro-scale robots consisting of light-reactive, contractile units. The ability to illicit controlled locomotion from optical illumination will open up new possibilities for achieving swarm-like behaviours in micro-scale systems.

Acknowledgements This work was supported by the EPSRC Robotics and Autonomous Systems Centre for Doctoral Training, FARSCOPE, UKRI Grant No. EP/S021795/1 (M. Uppington); and by the European Union under grant agreement 101070918, UKRI Grant No. 10038942 (S. Hauert.). This paper is an extension of our previous work 'Evolving morphologies for locomoting micro-scale robotic agents', MARSS © 2022 IEEE. Reprinted, with permission, from M. Uppington, P. Gobbo, S. Hauert and H. Hauser

Author Contributions M. Uppington wrote the majority of the main manuscript text with small inputs and suggestions from all authors. M. Uppington designed the methodology, implemented the simulation and artificial evolution trials, prepared all figures and supplementary materials, and wrote the analysis of the results gathered. All authors reviewed the manuscript

Funding This work was supported by the EPSRC Robotics and Autonomous Systems Centre for Doctoral Training, FARSCOPE, UKRI Grant No. EP/S021795/1 (M. Uppington); and by the European Union under grant agreement 101070918, UKRI Grant No. 10038942 (S. Hauert)

Data Availability Statement Our code can be accessed at <https://github.com/MattUppington/ProtocellSimulator>. Any queries should be directed to the first author, M. Uppington

Declarations

Competing interests The authors have no relevant financial or non-financial interests to disclose

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Negro A, Cherbuin T, Lutolf MP (2018) 3D Inkjet printing of complex, cell-laden hydrogel structures, scientific reports 8(1)
- Kriegman S, Blackiston D, Levin M, Bongard J (2020) A scalable pipeline for designing reconfigurable organisms, Proceedings of the National Academy of Sciences of the USA
- Wegner S, Sentürk OI, Spatz JP (2015) Photocleavable linker for the patterning of bioactive molecules'. Nature Publishing Group 5:2015
- Downs FG, Lunn DJ, Booth MJ, Sauer JB, Ramsay WJ, Klemperer RG, Hawker CJ, Bayley H (2020) Multi-responsive hydrogel structures from patterned droplet networks. Nat Chem
- Gobbo P, Patil AJ, Li M, Harniman R, Briscoe WH, Mann S (2018) Programmed assembly of synthetic protocells into thermoresponsive prototissues. Nat Mater 17(12)
- Galanti A, Moreno-Tortolero R, Azad R, Cross S, Davis S, Gobbo P (2021) A Floating Mold Technique for the Programmed Assembly of Protocells into Protocellular Materials Capable of Non-Equilibrium Biochemical Sensing. Adv Mater
- Denniss AMR, Gorochoowski TE, Hauert S (2019) Augmented reality for the engineering of collective behaviours in microsystems, in Proceedings of MARSS, IEEE Inc
- Denniss AR, Gorochoowski TE, Hauert S (2022) An Open Platform for High-Resolution Light-Based Control of Microscopic Collectives, Advanced Intelligent Systems
- Pfeifer R, Gomez G (2009) Morphological computation - Connecting brain, body, and environment. Lecture Notes in Computer Science. Springer, Berlin, Heidelberg, pp 66–83
- Slavkov I, Carrillo-Zapata D, Carranza N, Diego X, Jansson F, Kaandorp J, Hauert S, Sharpe J, (2018) Morphogenesis in robot swarms. Sci Robotics 3(25)
- Mautner C, Belew RK (2000) Evolving robot morphology and control. Artif Life Robotics 4(3):130–136
- Vujovic V, Rosendo A, Brodbeck L, Iida F (2017) Evolutionary developmental robotics: Improving morphology and control of physical robots. Artif Life 23(2)
- Miskin MZ, Cortese AJ, Dorsey K, Esposito EP, Reynolds MF, Liu Q, Cao M, Muller DA, McEuen PL, Co-hen I (2020) Electronically integrated, mass-manufactured, microscopicrobots. Nature. 584(7822):557–561
- Corteautes J, Egerstedt M (2017) Coordinated control of multi-robot systems. A Survey, SICE Journal of Control, Measurement, and System Integration 10(6):495–503
- Gautam A, Mohan S (2012) A review of research in multi-robot systems. IEEE 7th International Conference on Industrial and Information Systems (ICIIS), pp 1–5
- Xu K, Yang Y, Li B (2021) Brownian cargo capture in mazes via intelligent colloidal microrobot swarms. Adv Intell Syst 3(11)
- Uppington M, Gobbo P, Hauert S, Hauser H (2022) Evolving morphologies for locomoting micro-scale robotic agents , '2022 International Conference on Manipulation, Automation and Robotics at Small Scales (MARSS) pp 1–5. <https://doi.org/10.1109/MARSS55884.2022.9870459>
- Sitti M, Wiersma D (2020) Pros and Cons: magnetic versus optical microrobots. Adv Mater 32(20)
- Go G, Kwak D, Piao L, Choi H, Jeong S, Lee C, Park B, Ko S, Park J, Park S (2013) Manipulation of micro-particles using a magnetically actuated microrobot. Mechatronics 23(8)
- Hu W, Lum GZ, Mastrangeli M, Sitti M (2018) Small-scale soft-bodied robot with multimodal locomotion. Nature. 554:81–85
- Xu T, Zhang J, Salehizadeh M, Onaizah O, Diller E (2019) Millimeter-scale flexible robots with programmable three-dimensional magnetization and motions. Sci Robotics 4(29)
- Xie H, Sun M, Fan X, Lin Z, Chen W, Wang L, Dong L, He Q (2019) Reconfigurable magnetic microrobot swarm: Multimode transformation, locomotion, and manipulation. Sci Robotics 4:
- Lv C, Nambodiri Varanakkottu S, Baier T, Hardt S (2018) Controlling the trajectories of Nano/Micro particles using light-actuated marangoni flow. Am Chem Soc Nano Lett
- Zeng H, Wasylczyk P, Parmeggiani C, Martella D, Buresi M, Wiersma D (2015) Light-fueled microscopic walkers. Adv Mater 27(26)
- Sims K (1994) Evolving 3D morphology and behavior by competition. Artif Life 1(4)
- Gorochoowski TE, Matyjaszkiewicz A, Todd T, Oak N, Kowalska K, Reid S, Tsaneva-Atanasova KT, Savery NJ, Grierson CS, Di Bernardo M (2012) BSIm: An agent-based tool for modeling bacterial populations in systems and synthetic biology. PLoS ONE
- Y. Gu, X. Zhang, Q. Wu, Y. Li, B. Zhang, F. Gao, and Y. Luo (2019) Research on motion evolution of soft robot based on VoxCAD, in Lecture Notes in Computer Science. Springer Verlag
- Cickovski T, Aras K, Alber MS, Izaguirre JA, Swat M, Glazier JA, Merks RM, Glimm T, Hentschel HGE, Newman SA (2007) From genes to organisms via the cell a problem-solving environment for multicellular development. Comput Sci Eng 9(4):50–60
- Swat M, Thomas GL, Belmonte JM, Shirinifard A, Hmeljak D, Glazier JA (2012) Multi-Scale modeling of tissues using CompuCell3D. Comput Methods Cell Biol 110:
- Hirashima T, Rens EG, Merks RMH (2017) Cellular Potts modeling of complex multicellular behaviors in tissue morphogenesis. Development, Growth and Differentiation 59(5)
- Stokes G (1856) On the effect of the internal friction of fluids on the motion of pendulums. Cambridge Philosophical Society vol, Trans, p 9

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.