Joachimczak, M., Kowaliw, T., Doursat, R. & Wróbel, B. (2013) Controlling development and chemotaxis of soft-bodied multicellular animats with the same gene regulatory network. *Advances in Artificial Life: Proceedings of the 12th European Conference on the Synthesis and Simulation of Living Systems (ECAL 2013)*, September 2-6, 2013, Taormina, Italy; Liò, P., Miglino, O., Nicosia, G., Nolfi, S. & Pavone, M.: pp. 454-461. MIT Press, ISBN 978-0-2623-1709-2.

Controlling development and chemotaxis of soft-bodied multicellular animats with the same gene regulatory network

M. Joachimczak[1], T. Kowaliw[2], R. Doursat[2,3], and B. Wróbel[1,4]

[1]Systems Modeling Laboratory, IO PAN, Sopot, Poland
[2]Institut des Systèmes Complexes Paris Île-de-France (ISC-PIF), CNRS, Paris, France
[3]School of Biomedical Engineering, Drexel University, Philadelphia, USA
[4]Evolutionary Systems Laboratory, Uniwersytet im. Adama Mickiewicza, Poznań, Poland
mjoach@iopan.gda.pl

Abstract

The ability to actively forage for resources is one of the defining properties of animals, and can be seen as a form of minimal cognition. In this paper we model soft-bodied robots, or "animats", which are able to swim in a simulated two-dimensional fluid environment toward food particles emitting a diffusive chemical signal. Both the multicellular *development* and *behaviour* of the animats are controlled by a gene regulatory network (GRN), which is encoded in a linear genome. Coupled with the simulated physics, the activity of the GRN affects cell divisions and cell movements during development, as well as the expansion and contraction of filaments connecting the cells in the swimming adult body. The global motion that emerges from the dynamics of the animat relies on the spring-like filaments and drag forces created by the fluid. Our study shows that it is possible to evolve the animat's genome (through mutations, duplications and deletions) to achieve directional motion in this environment. It also suggests that a "minimally cognitive" behaviour of this kind can emerge without a central control or nervous system.

Introduction

In biological multicellular organisms, the dynamics of gene regulatory networks (GRNs) controls not only the growth of the organism, including the maintenance of the cells and overall structure, but also its behaviour. A striking example can be observed in social amoeba such as *Dictyostelium* (slime mold), where gene regulation controls both the aggregation of single cells into a slug and the adaptive behaviour of this collective entity (Bonner, 2008).

In theoretical biology and artificial life, artificial gene networks are used to understand how computational properties of biological networks evolve. One area of research is the evolution of control of multicellular development (Dellaert and Beer, 1996; Eggenberger Hotz, 1997; Doursat, 2009; Schramm and Sendhoff, 2011), another is computation in a more general sense (Banzhaf, 2003; Nicolau et al., 2010; Lopes and Costa, 2012). We addressed these two areas in our previous research using the artificial life system that we created, GReaNs (for Genetic Regulatory evolving artificial Networks; reviewed in Joachimczak and Wróbel, 2011). We investigated in particular the evolution of signal processing using continuous or spiking computational units (Joachimczak and Wróbel, 2010; Wróbel et al., 2012), and the evolution of soft-bodied artificial organisms, or "animats", whose development and locomotion were controlled by a GRN (Joachimczak and Wróbel, 2012; Joachimczak et al., 2012).

The use of a developmentally inspired stage to generate the morphology of a virtual robot is an active area of research, involving a range of abstractions for cellular and genetic control (Hornby and Pollack, 2002; Bongard and Pfeifer, 2003; Kowaliw et al., 2004; Doursat, 2008; Meng et al., 2011). The main contribution of our system lies in the combination of a biologically realistic encoding of the GRN (and genetic operators that allow for their complexification) with a realistic physics simulation. Physics rules govern the movement of cells during development, and the drag forces during locomotion in the fluid. Although in our current implementation the environment and the animats are two-dimensional, the system could be extended to 3D to make our results even more relevant. Physically plausible robots could take advantage of their softness—and thus resistance to damage and external forces—when interacting with other objects (for example, changing shape to squeeze through small openings). Although the properties of non-rigid, modular bodies have been explored before (Shimizu et al., 2005; Umedachi et al., 2010; Schramm and Sendhoff, 2011; Doursat et al., 2012; Hiller and Lipson, 2012; Rieffel et al., 2013), including our previous work on the diversity of locomotion strategies in soft-bodied animats (Joachimczak et al., 2012), the present study is the first attempt, to our knowledge, at evolving a fully decentralized controller and morphology of elastic animats that can sense and navigate their environment.

In the present paper we consider the evolution of gene regulatory networks able to control both the development of a soft-bodied animat and its emergent *multicellular chemotaxis*, a basic behaviour that consists of moving toward the source of an external signal. Despite its apparent simplicity, this task requires generating motion and coordinating numerous local cell actions to turn the body in the direction of a gradient. We identify and analyze here several morpholo-

gies and behavioural strategies toward this goal. The main contribution of this paper is to demonstrate that such *minimal cognition* (van Duijn et al., 2006) can collectively evolve in a multiagent system. We also suggest that it could be the first step toward more advanced cognitive abilities (Wróbel, 2012). Another novelty is a simplification of the artificial physics: instead of keeping a different set of environmental conditions for the developmental phase and the behavioural phase, we adopt the same physical rules for both.

Controlling the development and behaviour of multicellular soft-bodied animats

The model used in this paper builds upon our previous work on soft-bodied multicellular animats (Joachimczak et al., 2012; Joachimczak and Wróbel, 2012). As before, the gene regulatory networks that control the bodies are encoded in linear genomes. We provide here only a brief summary of how the encoding works and how the dynamics of the network is simulated, then we describe in more detail two modifications that we brought to the system: unifying the physical conditions of the developmental and behavorial phases, and designing new cell-to-cell communication. In this new design, chemical signals diffuse between cells under a constraint of conservation of their total amount (in the previous implementation, mass was not conserved).

Genome and gene regulatory network

The genome is represented by a list of genetic "elements" without fixed length. Genetic elements belong to three classes: (i) *genes*, which code for *products* (transcription factors or chemicals diffusing between cells), (ii) *regulatory elements*, and (iii) *special elements*, which encode inputs and outputs of the regulatory network. One or more regulatory elements form a regulatory region, which can be followed in the genome by one or more genes to make a *regulatory unit*. The activation levels of the regulatory elements of a unit influences the concentrations of products coded by this unit's genes (Fig. 1). Conversely, regulatory elements are activated by the products currently present in the cell, which virtually "bind" to the genome with various probabilities (related to their concentration) and various affinities to the regulatory sites.

To simulate the behaviour of a cell, we first decode the genome to obtain the corresponding GRN, in which nodes represent regulatory units, and weighted directed edges represent relations of regulation. The signs of the weights indicate whether the regulation is excitatory or inhibitory, while the weights tune the chemical affinity between products and regulatory elements. The affinity also depends on the "distance" between two elements, calculated by construing each element as a point in an abstract 2D space of chemical interactions (not to be confused with the physical 2D space of the animat). The affinity is set to 0 if the distance is above a certain threshold, and to a maximum value if the two points overlap.

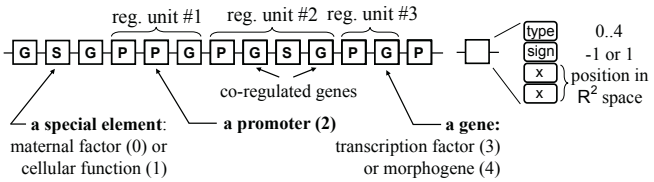

Figure 1: Genome (left) and structure of a genetic element (right). Each element consists of a *type* field, which specifies its class (G: gene, P: regulatory, S: special), a *sign* field, and N abstract *coordinates* (here, $N = 2$), which determine its affinity to other elements based on distance in $\mathbb{R}^N$.

While each cell in the animat body contains the same GRN, the product concentrations that encode the dynamic state of this network can be different from cell to cell. Concentrations are real values updated in discrete time steps. The increase in concentration, or "synthesis rate", of a product P is influenced by the concentrations of products that have a non-zero affinity to the regulatory elements of the unit encoding P. The combined effect of all the products binding to the same regulatory element is additive; the combined effect of all the regulatory elements in a regulatory unit is also additive. If the net effect of the products that have an affinity to the regulatory elements of a regulatory unit is negative, the products encoded by this unit's genes will decrease in concentration, or "degrade". In addition, the products encoded by regulatory units degrade spontaneously.

The minimal concentration of a product is always 0, but the maximum concentration is different for transcriptional factors (1.0) and diffusive products (10.0). There are two reasons for this 10-fold difference. First, if the maximum concentration of diffusive products was low, it could not be detected in the cells far away from the source cell. Second, when the initial population is formed during simulated evolution (i.e. when the genomes are constructed randomly for the individuals in this population), elements that code for diffusive products are introduced in these genomes less often than elements that code for transcription factors. In contrast to our previous model, the products diffuse here in the body along the filaments that connect the cells (both during development and locomotion). At each time step, the fraction of concentration of a diffusive product transferred between cells is proportional to the difference of concentrations between these cells.

Diffusive products can be considered to be one form of an output produced by a cell (and input received by other cells). Our genome model also includes elements that encode GRN inputs coming from the environment and outputs representing cell actions. These *special elements* are not tied to regulatory units, the graph nodes to which they correspond do not have recurrent connections, and direct connections between input and output nodes are not allowed.

Input elements behave like other regulatory products (transcription factors and morphogens), but their concentration represents an environmental signal. In this paper we use five types of input elements, four of which can be seen as encoding "maternal morphogens". Three of these morphogens diffuse during development from three point sources, so their perceived concentration in a given cell depends on the current position of this cell. One environmental signal is always present in the same concentration (1.0), throughout development and beyond, when the animat moves (this signal plays the same role as a bias node in artificial neural networks). The fifth chemical starts diffusing in the environment when the animat has finished developing, at which point the animat is supposed to move toward the source of this chemical. Its concentration in each cell depends on the cell's distance from the source, and goes to 0 for a distance larger than 400 units (noting that the expected value from multiple trials of the initial distance from the center of mass of the animat to the food source is 300).

Whereas input elements encode products whose concentration is determined by the environment, output elements encode products whose concentration impacts the behaviour of the cells and the entire animat after development. In this study we use five output elements representing five possible cell actions: (i) division (when the concentration of the corresponding product crosses a threshold), (ii) rotation to the left and (iii) rotation to the right after division (cell orientation is represented by a vector; the rotation angle depends on the concentration of two products), (iv) contraction and (v) expansion of the filaments linked to the cell (the two products corresponding to these last actions are used only after development, when the mature body is able to move).

Physics of cell interactions

As in our previous work, the animats are spring-mass systems in which cells correspond to point masses, and neighbouring cells are connected by filaments that act as weightless springs. This neighbourhood relation is determined by calculating the Gabriel graph (Gabriel and Sokal, 1969) of cell positions (Fig. 2).

In our model of two-dimensional swimming, taking after the simulation of undulatory robotic locomotion by Sfakiotakis and Tsakiris (2006), the fluid is stationary and only the spring-edges on the outline of the animat are subject to fluid drag. The force exerted on an edge of length L is the sum of a tangential component $F_T = -d_T L\, v_T^2 \,\mathrm{sign}(v_T)$ and a normal component $F_N = -d_N L\, v_N^2 \,\mathrm{sign}(v_N)$, both proportional to the squares of the respective velocity components v_T and v_N via fluid drag coefficients d_N and d_T (where $d_N = 200\, d_T$).

Soft bodies during development and locomotion In contrast to our previous work, the rules of physics governing the development and locomotion are identical. These two phases remain separate, however, and are different in three respects. First, to prevent excessive forces and movements in the developing embryo due to cell division, cells are slowed down by an extra drag component proportional to the square of their velocity. This correction can be interpreted as the presence of an intracellular fluid more viscous than the external fluid (an alternative, not used here, would be to consider that immature filaments are less stiff). The second difference is that mature filaments in the adult body define polygons that act as pressurized chambers, whose expansion and contraction drive them out of equilibrium and generate a pressure force along the normal of each edge. Thus these chambers constitute a "hydrostatic skeleton" for the animat, which also prevents cells from passing through filaments. Finally, during development cells can break connections or form new ones (as if sprouting filaments or destroying them), whereas the connectivity in the locomoting adult remains fixed.

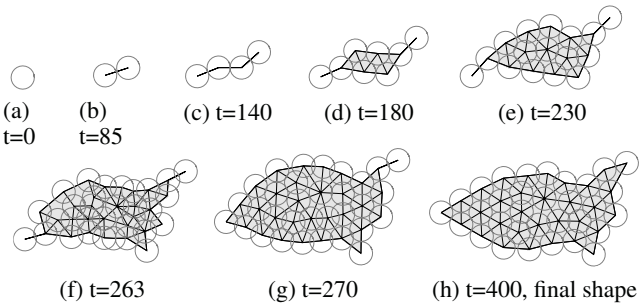

Figure 2: Example of the developmental mechanics (for individual #1, shown in Fig. 3). Cells are represented as circles of radius r and connected by springs with resting length $2r$.

Formally, this means that the Gabriel graph is recalculated at every step of the development, and each pair of neighbouring cells is connected by a spring whose resting length is the sum of the cells' radii (in the experiments described here, all cells have the same radius). When a cell produces a daughter cell through division, the new cell is placed closer than the sum of the radii, so the spring that connects these two cells pushes them away from each other, creating a cascading effect in the body. As the organism is growing, cells always attempt to maintain constant distances between them (Fig. 2), and new neighbourhood relations lead to the creation or removal of springs. To keep computational costs reasonable, we used a hard limit of 32 cells.

Once the development is finished, the filaments "mature" i.e. although they retain their elasticity and the body may change shape during movement, the pattern of connections between cells is no longer modified. The initial resting length L_0 of each spring is set to the length it had at the last time step of development. From this point, each cell controls the springs connected to it using the products encoded by two output elements: one product for expansion and one for contraction. The concentrations of expansion

products e_1 and e_2 in the two cells connected by a spring, and the concentrations of contraction products c_1 and c_2 combine additively to modify the resting length according to $L = (1 + A_{\max}(e_1 + e_2 - c_1 - c_2))L_0$, where $A_{\max}$ is a global parameter of the system representing the maximum actuation amplitude ($A_{\max} = 0.2$ in this paper).

Evaluation of behaviour and chemotaxis

In our evolutionary model, genetic operators can add elements (duplications), remove elements (deletions), or change the elements' type, sign, and coordinates (point mutations). The first two operations affect the size of the genome and the number of nodes and edges in the GRN, while the change of coordinates affects the affinities between products and regulatory elements. The genetic algorithm is generational, with population size 100, and tournament selection on five randomly drawn individuals. Five of the individuals in each generation are propagated without change to the next generation (elitism), and 20% undergo sexual reproduction (multipoint crossover). An evolutionary run stops when the fitness value is stable over a 500-generation span, which happened between generations #2000 and #3000 in the experiments described here. To accelerate evolution and evaluation, nonviable individuals are removed from the population, where an individual is deemed "viable" if three conditions are met: (i) there is a path between at least one input and the outputs associated with division, contractions or expansions in the GRN, (ii) no cell division happens during the last 100 time steps before the end of development (to allow the physics to equilibrate the adult structure; there is a fixed number of time steps for development), and (iii) the concentrations of expansion or contraction products vary during locomotion. These criteria of viability guide the search for the random genome that will be used to create the genomes of all the individuals in the initial population. These individuals are generated from the seed genome via the duplication, deletion and point mutation operators. The random search of a seed genome requires a few thousand trials.

Fitness evaluation

After the soft body has fully developed, through cell divisions starting from a single cell, the animat begins to move. In our preliminary experiments we placed a food particle repeatedly at eight random locations forming a circle around the animat's center of mass, and gave higher fitness to animats closer on average to the particle (after a fixed number of time steps). Yet, the evolutionary search in this scenario was not very efficient: only about one third of the evolutionary runs resulted in "champions" that showed some chemotaxis abilities, but at a considerable computational cost due to the required eight test cases for each individual in each generation.

To improve the efficiency of the evolutionary search, we redesigned the fitness function to be composed of five terms obtained by evaluating an individual in five test situations. (1) The first test situation assessed the ability to move as such: we measured the distance travelled in 10,000 simulation steps. This evaluation stage also allowed to determine the main axis of the animat and its preferred direction of movement. (2) Then, we placed a food particle on the animat's left, between -30° and -90° from the main axis, at a distance chosen uniformly and randomly in the range [200,400] from its center of mass (animats cover about 100 units along the main axis), and measured the remaining distance to the particle after 15,000 simulation steps or, if the animat's body overlapped with the particle earlier, the time it took. (3) The third test repeated the second: the state of the animat including its shape and the concentrations of products was reset to the state it had at the end of the first test, and the particle was placed again on the left. (4, 5) The last two tests were similar to tests (2, 3) with the particle placed on the right. The resulting fitness function (maximized by the genetic algorithm) was a linear combination of the distance d travelled in the first test (via a increasing reward), the remaining distances d_n from the animat's center of mass to the food particle, and the total durations t_n of the last four tests (via decreasing rewards):

$$f_{\text{fit}} = \frac{d}{c_m} + \sum_{n=1}^{4} \left(1 - \frac{d_n}{c_f} + s_r\left(1 - \frac{t_n}{t_{\max}}\right)\right), \quad (1)$$

where c_f is the maximum distance at which a particle could be placed, $t_{\max} = 15,000$ is the maximum number of steps in each test, s_r is the weight of the time reward with respect to the remaining distance reward (here, $s_r = 4$), and $1/c_m$ is the weight of the distance travelled with respect to the last four tests. This coefficient was set to a value such that, for an individual with efficient locomotion and chemotaxis, the first reward component was of the same order as each of the other four reward components. Our fitness function design promotes the evolution of a simpler behaviour first (here, locomotion), so that a more complex one (chemotaxis) can be built upon it. Considering the relations between learning and evolution, this design brings the fitness function close to a trainer or tutor that promotes gradual development of competences (by "scaffolding" or "shaping" the agent; (Wood et al., 1976; Dorigo and Colombetti, 1994)).

Results: swimming patterns of four champions from independent evolutionary runs

An analysis of the champions obtained from multiple independent evolutionary runs ($n = 40$) shows that about half of them were able to change direction and to head toward the food particle, while the other half could only swim forward. Our previous work (Joachimczak et al., 2012) had identified four classes of morphologies and styles of motor behavior that emerged more distinctly among the continuum of possible scenarios: symmetrical protrusions on the left

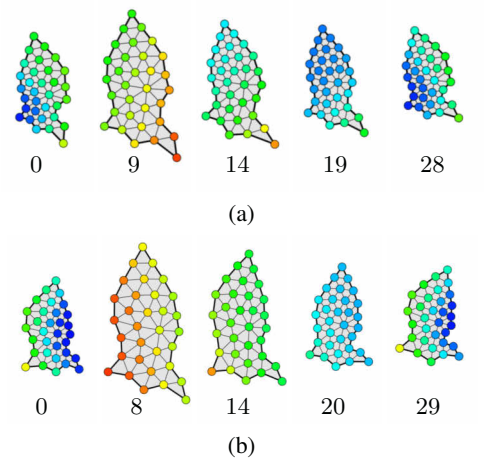

Figure 3: Patterns of cell activation (i.e. concentration of expansion minus contraction products) in individual #1 during one motion cycle while it performs a left (a) or right (b) turn. Red indicates expansion (positive activation); green: resting length; blue: contraction (negative activation). Numbers indicate time steps. The animat swims upward. Videos available at: `http://youtu.be/zi3p164aefY` and `http://youtu.be/Cqt8Fy3CWlA`

and right (or "fins"), a protrusion at the end (or "tail"), undulation of the whole body, and alternation of whole-body pulsations, consisting of either fast expansion and slow contraction of bodies that had a pointy front and a blunt end, or the other way around (fast contraction with a blunt front). Strategies based on pulsations worked by exploiting the fact that fluid drag is proportional to velocity squared. It was also characterized by rapid swings of the concentration of expansion and contraction products in all cells at the same time, whereas in the first three strategies these concentrations varied in a sinusoidal fashion and exhibited phase gradients along the axes of the body. In the present work, it is the pulsation strategy that happened to be the most common among the champions who showed efficient chemotaxis—despite the fact that, in our previous work, individuals with protrusions were the fastest in forward motion. It is interesting to note that the other three strategies also tend to appear in the experiments reported here, but less clearly or only partially as components of a mix (see examples below). This is probably due to the different physics model and the new requirements for chemotactic abilities which could be encouraging pulsations over protrusions or undulation.

We chose four animats among the fittest to be analyzed in greater detail. The pulsation strategy is used by the first three, among which two have elongated bodies in the direction of motion. Individual #1 exhibits a sharp front and a blunt end, contracts slowly and expands quickly (Fig. 3). Individual #2 shows the opposite, with a blunt front and sharp end, contracting quickly and expanding slowly. It also generates thrust by wiggling a "tail" (Fig.4). Individual #3 also uses a mixed strategy, generating thrust in part from a pul-

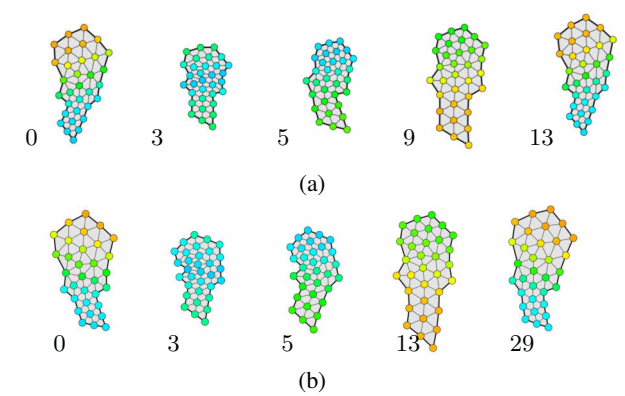

Figure 4: Patterns of cell activation (i.e. concentration of expansion minus contraction products) in individual #2 during one motion cycle while it performs a left (a) or right (b) turn. Red indicates expansion (positive activation); green: resting length; blue: contraction (negative activation). Numbers indicate time steps. The animat swims upward. Videos available at: `http://youtu.be/TS8Q0JfI7o0` and `http://youtu.be/Dw8-YCWodn8`

sating clump in the middle (by fast expansion and slow contraction) and in part by the movement of a small tail, too. A wave of contraction travelling from the front to the back moves this tail in a position perpendicular to the main axis of the animat (equal to the direction of motion) when the animat's back expands, so that the tail pushes the animat forward (Fig. 5). Individual #4 is sharply different from the other three, as its body is elongated in the direction perpendicular to the main axis, and it moves by using two joined "fins", which push backward in synchrony to generate a forward movement (Fig. 6). These fins expand when moving backward, and contract on their return. Their motion is induced by a wave of contractions travelling from the back toward the front.

In all four animats the control of chemotaxis performed correctly in the more general situation in which, after the animat reached one food particle, we placed another particle away from the animat without resetting its state to a pre-food situation (although the state was reset during the evaluation phase of the genetic algorithm). When turning toward the food, these four animats did not change their motion pattern, thus it is not immediately obvious how they performed the turn. Defining the level of "activation" of a cell to be the concentration difference between the expansion product and the contraction product, changes in collective activation patterns between the left and right turn can be observed in individuals #1 and #4, while this symmetry is much less clear in individuals #2 and #3 (Figs. 3-6).

To understand how the control of turning worked, we compared the average activation of each cell when food was placed in front of the animat to the activation of each cell when food was placed on the left or on the right. The experiments were performed as follows: (i) the animat was

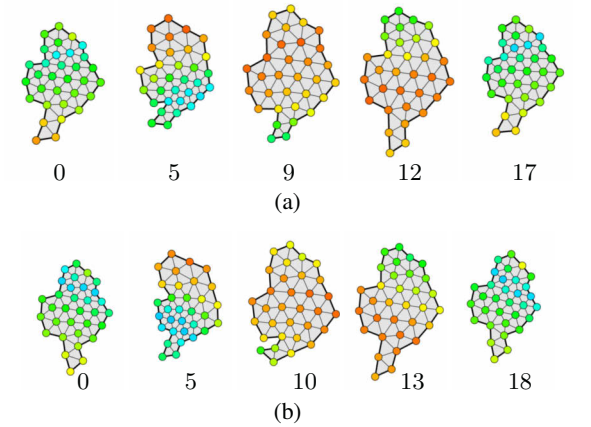

Figure 5: Patterns of cell activation (i.e. concentration of expansion minus contraction products) in individual #3 during one motion cycle while it performs a left (a) or right (b) turn. Red indicates expansion (positive activation); green: resting length; blue: contraction (negative activation). Numbers indicate time steps. The animat swims upward. Videos available at: `http://youtu.be/-HoN7ZGU6W4` and `http://youtu.be/UZeCkWgeA5Q`

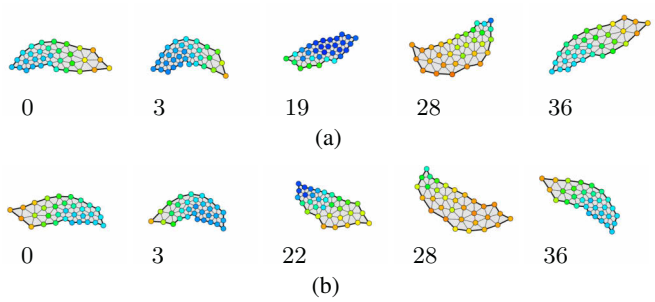

Figure 6: Patterns of cell activation (i.e. concentration of expansion minus contraction products) in individual #4 during one motion cycle while it performs a left (a) or right (b) turn. Red indicates expansion (positive activation); green: resting length; blue: contraction (negative activation). Numbers indicate time steps. The animat swims upward. Video available at: `http://youtu.be/rvM2T8gpDnU`

allowed to move forward without a particle for 10,000 steps and the direction of movement in the last 50 simulation steps was used to determine its main axis, then the average cell activity was calculated (over 5,000 steps in each case) with (ii) a food particle placed in the front, (iii) 60° to the left, and (iv) 60° to the right. The initial state of the animat (the product concentrations and body shape) in each case was identical to the state at the end of step (i).

Analysis of the changes in the average cell activation indicates that pulsating animats use different strategies based on changing the size of their body parts. Individual #1, the animat that contracts slowly and expands quickly, turns by contracting the part of its body closest to the food (Fig. 7a), as does individual #3, which uses a mixed strategy of pulsation and tail propelled by a wave of contractions (Fig. 7c). On the other hand, individual #2, which contracts quickly and expands slowly, turns by expanding the side opposite to food (Fig. 7b).

In terms of how forces generate motion, when individual #1 with a blunt end expands quickly, it pushes the blunt end against the fluid, and the relative increase in length of the external edges on the right side causes a push toward the left. On the contrary, when individual #2 with a pointy end contracts quickly, a relative increase in edge length on its right side leads to an additional pull toward the left. Although individuals #1-3 showed similar strategies for controlling turns, only in the case of the individual #1, which used a pure pulsation strategy, did we observe an immediate reaction of the cells (i.e. contraction of the springs, but maintaining their pulsation) to a food particle close to them. We also observed similar contractions for cells close to the particle when it was placed inside the body, although this was not experienced during evolution.

In comparison to the first three examples of animats, all propelled by pulsations, individual #4 using two joined fins can turn and move significantly faster, sometimes even overshooting the target but correcting its trajectory afterwards. It is able to switch the direction of the wave of contractions: without food, the wave moves from the back to the front; with a food particle on the right, it moves from the right tip to the left tip, and vice-versa (Fig. 7d). During the switch, the animat maintains the overall motion pattern: the contraction waves are synchronized so that when the right half of the body moves backward, this part contracts. This results in a lower thrust from the right fin, hence a right turn, which is captured by the analysis of average cell activation.

Because the chemical diffusing from the food particle is sensed by all the cells, it is conceivable that the gene network is reacting proportionally to the strength of the incoming signal, and using this direct response to stimulate the turn. To detect this possibility, we computed the Spearman's rank correlation coefficient between the activity (expansion and contraction) of the cells and the distance to the food, i.e. strength of the diffusive signal (Fig. 8). The largest correlations can be observed in individual #1, the pulsating animat that expands quickly. During this animat's behaviour, the high correlations vary in regular patterns. This indicates that cells close to the food expand their springs and cells far from food contract their springs proportionally to the food signal. A similar pattern can often be seen, although not as clearly, in the other two pulsating animats, #2 and #3. This is not, however, the case for the two-finned individual #4. In its case, correlations between activity and distance are relatively low and unstructured, which suggests a more complex strategy during behaviour, including a substantial role played by intercellular communication.

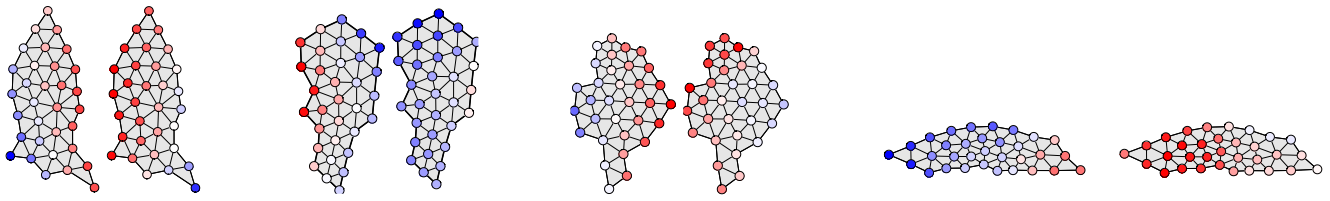

(a) (b) (c) (d)

Figure 7: Change in average cell activation (i.e. concentration of expansion minus contraction products) when moving a food particle from the left or right side to the front of an individual. White cell: no change; blue: average activation is lower; red: higher. The four pairs from (a) to (d) correspond respectively to individuals #1 to #4. All animats shown moving upward.

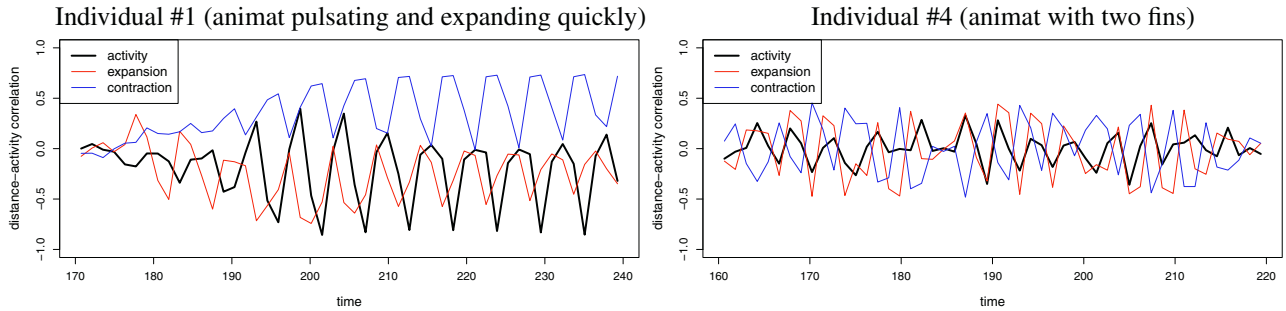

Figure 8: Variations of the Spearman correlation between cell activity and distance to food. Three curves are shown for individuals #1 and #4. They represent the correlation between cell activity X and the cells' absolute distance to food, where X stands for the cells' expansion (in red), their contraction (in blue), and their activity (expansion minus contraction, in black). The total period (60 to 70 time steps) corresponds to an animat turning left and moving toward the food. Values near 0 mean no correlation; values near $+1$ and -1 mean high correlations, which indicate here a direct relationship between the cells' behaviour and their distance to the food.

Summary, perspective, and future work

In this work, we showed that it was possible for an evolutionary process to successfully produce soft-bodied multicellular animats that can forage for resources in their environment, and whose embryonic development and adult behaviour are both controlled by the same gene regulatory network. Neither a particular morphology nor a particular type of control were enforced, but different strategies were discovered by evolution starting from random genomes. The chemotactic behaviour that emerged from the interplay between the shape of the body and the local responses of differentiated cells, in the absence of any central control, can also be regarded as "minimally cognitive".

The simulated evolution of a coordinated collective behaviour, where multiple agents are driven by the same distributed controller, provides a way to explore the space of possible morphologies and efficient modes of control in the nascent field of soft-bodied robotics. Our results show that soft robots are able to navigate efficiently and robustly by pulsing the body (symmetrically along the main axis) and expanding their left or right side slightly more in order to turn. We hope that these results can contribute to providing a source of inspiration for the development of new materials and actuators for soft robots.

Before such materials and actuators are available, however, much work can still be accomplished in virtual environments. As future work, we plan to investigate more thoroughly the types of motion and the nature of communication and synchronization among the cells of evolved individuals. In particular, we want to analyze how intercellular communication works to achieve efficient behaviour by exploring a scenario where only a subset of cells (for example, the cells on the surface) can sense the environment, while the rest of the body must rely on indirect information passed through diffusive substances. We would also like to better assess the benefits of distributed control, especially in terms of robustness to damage (e.g. malfunction of springs or cells) and resistance to external disturbances (e.g. distractors and noise).

Acknowledgments

This work was supported by a PAN-CNRS collaborative project, IO PAN (task IV.3), the Polish National Science Centre (project BIOMERGE, 2011/03/B/ST6/00399), the Tri-city Academic Computer Center (TASK), and the Interdisciplinary Center for Molecular and Mathematical Modeling (ICM, University of Warsaw; project G33-8).

References

Banzhaf, W. (2003). On the dynamics of an artificial regulatory network. In *Advances in Artificial Life: Proc. of the 7th European Conference on Artificial Life (ECAL 2003)*, volume 2801 of *LNAI*, pages 217–227. Springer.

Bongard, J. C. and Pfeifer, R. (2003). Evolving complete agents using artificial ontogeny. In Hara, F. and Pfeifer, R., editors, *Morpho-functional Machines: The New Species*, pages 237–258. Springer Japan.

Bonner, J. T. (2008). *The Social Amoebae: The Biology of Cellular Slime Molds*. Princeton University Press.

Dellaert, F. and Beer, A. D. (1996). A developmental model for the evolution of complete autonomous agents. In *From Animals to Animats 4: Proc. of the 4th International Conference on Simulation of Adaptive Behavior (SAB 1996)*, pages 393–401. MIT Press.

Dorigo, M. and Colombetti, M. (1994). Robot shaping: developing autonomous agents through learning. *Artif. Intell.*, 71(2):321–370.

Doursat, R. (2008). Organically grown architectures: creating decentralized, autonomous systems by embryomorphic engineering. In Würtz, R. P., editor, *Organic computing*, Understanding complex systems, pages 167–199. Springer.

Doursat, R. (2009). Facilitating evolutionary innovation by developmental modularity and variability. In *Proc. of the 11th Annual Conference on Genetic and Evolutionary computation*, GECCO '09, pages 683–690. ACM.

Doursat, R., Sanchez, C., Dordea, R., Fourquet, D., and Kowaliw, T. (2012). Embryomorphic engineering: Emergent innovation through evolutionary development. In Doursat, R., Sayama, H., and Michel, O., editors, *Morphogenetic Engineering: Toward Programmable Complex Systems*, pages 275–311. Springer-Verlag.

Eggenberger Hotz, P. (1997). Evolving morphologies of simulated 3D organisms based on differential gene expression. In *Proc. of the 4th European Conference on Artificial Life (ECAL 1997)*, pages 205–213. MIT Press.

Gabriel, K. R. and Sokal, R. R. (1969). A new statistical approach to geographic variation analysis. *Syst. Zool.*, 18(3):259–278.

Hiller, J. and Lipson, H. (2012). Automatic design and manufacture of soft robots. *IEEE Trans. Robot.*, 28(2):457–466.

Hornby, G. S. and Pollack, J. B. (2002). Creating high-level components with a generative representation for body-brain evolution. *Artif. Life*, 8(3):223–246.

Joachimczak, M., Kowaliw, T., Doursat, R., and Wróbel, B. (2012). Brainless bodies: Controlling the development and behavior of multicellular animats by gene regulation and diffusive signals. In *Artificial Life XIII: Proc. of the 13th International Conference on the Simulation and Synthesis of Living Systems*, pages 349–356. MIT Press.

Joachimczak, M. and Wróbel, B. (2010). Processing signals with evolving artificial gene regulatory networks. In *Artificial Life XII: Proc. of the 12th International Conference on the Simulation and Synthesis of Living Systems*, pages 203–210. MIT Press.

Joachimczak, M. and Wróbel, B. (2011). Evolving gene regulatory networks for control of artificial cells: signal processing, chemotaxis, multicellular development. In *Proc. of the SynBioCCC Workshop*, pages 4–6.

Joachimczak, M. and Wróbel, B. (2012). Co-evolution of morphology and control of soft-bodied multicellular animats. In *Proc. of the 14th International Conference on Genetic and Evolutionary Computation*, GECCO '12, pages 561–568. ACM.

Kowaliw, T., Grogono, P., and Kharma, N. (2004). Bluenome: A novel developmental model of artificial morphogenesis. In *Conference on Genetic and Evolutionary Computation*, GECCO '04, pages 93–104.

Lopes, R. and Costa, E. (2012). The regulatory network computational device. *Genet. Program. Evol. M.*, 13(3):339–375.

Meng, Y., Zhang, Y., and Jin, Y. (2011). Autonomous self-reconfiguration of modular robots by evolving a hierarchical mechanochemical model. *IEEE Comput. Intell. Mag.*, 6(1):43–54.

Nicolau, M., Schoenauer, M., and Banzhaf, W. (2010). Evolving genes to balance a pole. In *EuroGP: 13th European Conference on Genetic Programming*, volume 6021 of *LNCS*, pages 196–207. Springer.

Rieffel, J., Knox, D., Smith, S., and Trimmer, B. (2013). Growing and evolving soft robots. *Artif. Life*, pages 1–20.

Schramm, L. and Sendhoff, B. (2011). An animat's cell doctrine. In *ECAL 2011: Proc. of the 11th European Conference on the Synthesis and Simulation of Living Systems*, pages 739–746. MIT Press.

Sfakiotakis, M. and Tsakiris, D. P. (2006). Simuun : A simulation environment for undulatory locomotion. *Int. J. Model. Simul.*, 26:350–358.

Shimizu, M., Ishiguro, A., and Kawakatsu, T. (2005). A modular robot that exploits a spontaneous connectivity control mechanism. In *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1899–1904. IEEE.

Umedachi, T., Takeda, K., Nakagaki, T., Kobayashi, R., and Ishiguro, A. (2010). Fully decentralized control of a soft-bodied robot inspired by true slime mold. *Biol. Cybern.*, 102(3):261–269.

van Duijn, M., Keijzer, F., and Franken, D. (2006). Principles of minimal cognition: Casting cognition as sensorimotor coordination. *Adapt. Behav.*, 14(2):157–170.

Wood, D., Bruner, J. S., and Ross, G. (1976). The role of tutoring in problem solving. *J. Child Psychol. Psyc.*, 17(2):89–100.

Wróbel, B. (2012). Challenges for a-life approach to artificial cognition: in search for hierarchy of cognitive systems. In *Artificial Life XIII: Proc. of the 13th International Conference on the Simulation and Synthesis of Living Systems*, pages 599–600. MIT Press.

Wróbel, B., Abdelmotaleb, A., and Joachimczak, M. (2012). Evolving spiking neural networks in the greans (gene regulatory evolving artificial networks) platform. In *EvoNet2012: Evolving Networks, from Systems/Synthetic Biology to Computational Neuroscience Workshop at Artificial Life XIII*, pages 19–22.