

Lateral gene transfer as a form of ultrastability

On the cybernetics of bdelloid rotifers

Pietro Galliani

September 19, 2016

1 Introduction

1.1 Reproduction and lateral gene transfer in Bdelloidea

Bdelloid rotifers are small (typically between 150 and 700 μm long [11]), worm-like invertebrates which are commonly found in freshwater habitats all over the world. Like all members of the *Rotifera* phylum, bdelloids feed on organic detritus, algae, bacteria and protozoa by mean of their *corona*, a ciliated structure whose movements cause water currents to bring food into their mouth. Some species are sessile, but most are capable of moving, either through free swimming or by crawling along a substrate.¹

One highly unusual aspect of the bdelloid lifecycle lies in the fact that no male has been observed in any species of this class, all of which reproduce exclusively by parthenogenesis. Other parthogenetic species of animals are of course known; however, whereas they generally tend not to speciate much and, often, to go extinct within a few thousands of years as they prove themselves incapable of adapting to environmental changes [14], the evolutionary record shows that bdelloids have enjoyed a respectable amount of evolutionary success for at least 35 to 40 millions of years (the age of oldest amber in which they have been identified [12]), evolving to adapt to different environments [13] and diversifying into more than 450 different species[5]; and for this reason, they have been described as “something of an evolutionary scandal” [9] and as “ancient asexual scandals” [8].

A possible explanation for this anomaly has been recently proposed in terms of massive amounts of lateral gene transfer [6, 3]. In brief, under stressful environmental conditions bdelloids can enter anhydrobiosis, rapidly dehydrate, and wait for up to nine years [7] for more favorable conditions to return; and as they recover and reconstitute their body, the process that they employ to repair possible damage to their DNA can incorporate foreign genes from their external environment, thus allowing for the development of genetic diversity despite their lack of sexual reproduction.

1.2 Homeostasis and ultrastability in the bdelloid life cycle

According to Ashby [1, 2], the behaviour of living organisms — up to and including abstract, high-level properties such as a sense of self, an aesthetic appreciation for microscopic creepy-crawlies, or a damnable tendency to procrastinate — is ultimately the outcome of *homeostatic activity* directed towards the maintenance of the values of certain *essential variables* within some boundaries. Through its interaction with a mutable world, this equilibrium-seeking activity leads the dynamics of the organism into a (possibly very rich) dance of adaptations and counter-adaptations, which in the can even lead to the organism “breaking” and radically altering its internal organization. As a proof of concept, Ashby famously created the *Homeostat*, an electromagnetic device whose subunits react to environmental changes by randomly changing their internal configuration (and possibly affecting the other subunits in turn) in order to keep certain values within their given limits. More recently, Di Paolo [4] investigated the effect of homeostasis on the adaptation of a phototaxic robot to visual inversion. After evolving the robot’s simple (but capable of plasticity) neural

¹For a video illustrating rotifer locomotion, see <https://www.youtube.com/watch?v=hed0cCfCcYQ>.



Figure 1: A bdelloid rotifer (uncertain species). Photo by Damián H. Zanette, released in the public domain by its author. Cropped photo, original available at http://en.wikipedia.org/wiki/File:Bdelloidea1_w.jpg

network in order to perform phototaxis and maintain the firing rates of its neurons within certain limits, Di Paolo swapped the inputs of its left and right sensors; and, after a first period of adaptation, the robot spontaneously reacquired its phototactic ability.

Even though Ashby focused on behavioural (as opposed to genetic) adaptations, it is hard not to see a connection between his notion of ultrastability — of an organism “breaking” in order to reconfigure itself and maintain some essential variables within their limits — and the bdelloids entering stasis in order to survive a stressful condition and, at the same time, acquire extraneous and novel genetic material. In both cases, environmental feedback can lead to a radical reconfiguration of the organism which, if successful, may better equip it (and, in the case of rotifers, its descendants) to preserve itself in a mutable and often unpredictable environment.

In what follows, I will explore this analogy through the simulation of a population of simple Braitenberg-like machines and its evolution under a few different scenarios, incorporating sexual/asexual reproduction, bdelloid-like gene transfer (or a lack of it), limited/unlimited lifespan ², and — as another possible mechanism for homeostasis and adaptation — increased (or non-increased) perceptual noise as a result of dangerous internal statuses.³ As we will see, lateral gene transfer will prove itself an excellent mechanism for adaptation, more efficient overall than damage-related noise. Despite its conceptual and mechanical similarity, on the other hand, sexual reproduction (in the form and for the parameters employed in this work) will show itself to be actively counterproductive to the development of homeostatic behaviour.

This confirms that — under some, evidently not all, conditions — bdelloid-like lateral gene transfer may be a viable mechanism for speciation and adaptation to changing environmental conditions, and sexual reproduction may not.

²Differently from other simple organisms, bdelloids are not as a rule biologically immortal — see for example [10] — but it seemed a reasonable enough variation to test. Also, I may have performed a good number of tests before noticing that I had allowed for unlimited lifespan.

³I did not attempt in any way to simulate the bdelloids’ quiescence period, however, as that did not appear particularly relevant to the topic at hand.



Figure 2: Ashby’s homeostat. Image taken from the slides of A. Seth’s course on Adaptive Systems (University of Sussex, 2014-2015).

2 Methods

2.1 The robots

The physical structure of the robots is, for the most part, similar to the one described in [4]. A robot’s body is circular, with a diameter of 10 units, two diametrically opposed motors and two sensors situated at an offset of 60 degrees (plus or minus 15 degrees⁴) from the direction of movement. As in Di Paolo’s work, the robots are assumed to be inertia-less, the sensors perceive light with an intensity proportional at the inverse square of the distance from the sources of light, and the sensors cannot see through the robot’s own body.⁵

In addition, and differently from [4], each sensor can distinguish between lights of two different “colors”, “red” and “blue”; and each robot has an internal reserve of “energy”, which they expend at a constant rate and recharge at a rate decreasing with their squared distance from the light sources. In order to survive, the robot must keep the value of this reserve between 0 and 110: below zero, it will die of “hunger”, whereas above 110 it will die of “overheating”. Normally, the perceptions only receive a small amount of noise from the sensors, chosen so to encourage the development of both exploratory and directed behaviours; but if the robot is within ten energy points from dying, either from hunger or from overheating, this noise amount can be made to increase tenfold. The robots do *not* have otherwise access to the value of their amount of stored energy: therefore, in order to perform homeostasis, they have to either “hard-code” within their genome the preferred distance from each light source or rely on this increased noise as an alarm signal (if available).

The robots are controlled by a simple ten-neuron continuous-time neural network. The neurons’ behaviour is described by the equations

$$\dot{y}_i = -y_i + \sum_j w_{ij} z_j + I_i; \quad z_j = \frac{1}{1 + \exp[-(y_j + b_j)]}$$

where y_i represents the membrane potential of the i -th neuron, z_i its firing rate, I_i its external sensory input (if any), b_i its bias and w_{ij} the weight of the connection from neuron i to neuron j . As seen from the above equations, and differently from Di Paolo’s work, the time constant has been fixed to one for all neurons and

⁴Di Paolo’s model allowed only for a variation of up to ± 5 degrees in eye angle; but since in this work I am studying the efficacy of bdelloid-style horizontal gene transfer, I wanted to prevent the results from relying too much on physical similarities between the robots. On the same note, and for the same reason, the variation in the eye position is *not* genetic, but is selected randomly at every robot birth.

⁵However, the robots *can* see through the bodies of the other robots participating to the simulation, as interaction between robots is not modeled in any way.

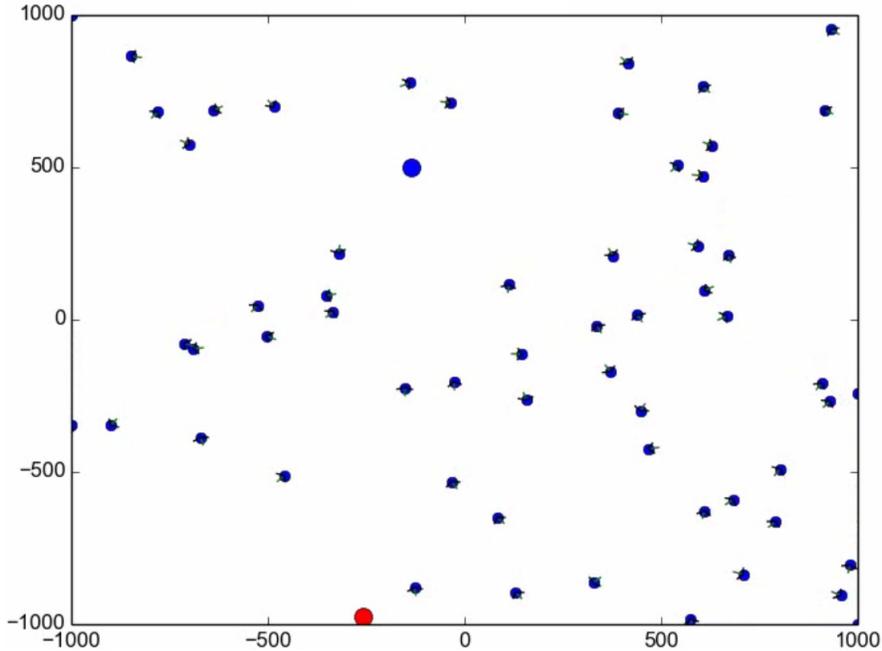


Figure 3: The arena: red light source, blue light source, and sixty randomly positioned robots.

all robots: indeed, allowing it to vary (from neuron to neuron, and from robot to robot) and encoding it in the robots' very genotype did not seem to lead to any novel form of behaviour, but merely slowed the speed of evolution. As in [4], the connection weights are allowed to take values⁶ between -8 and $+8$ and the biases are allowed to take values from -3 and $+3$; but differently from the above-mentioned paper, I did not allow for any form of synaptic plasticity.

Of the ten neurons, the first four receive as input the amounts of red and blue light received from each sensor, while the firing rates of the fifth and sixth one drive the activity of the motors. The remaining four neurons receive no direct external input and have no direct effect on the motors.

In some simulations, the robot lifespan has been limited to 500 seconds: and as we will see, some results suggests that this might have a significant effect on the evolution of the system, but further data should be gathered in order to confirm or deny this.

2.2 The environment

The robots are situated in a 1000×1000 bidimensional arena containing, at every time, two light sources — a red one and a blue one. The amount of energy released by either can be different, and can be changed by the user; but the amount of *light* released, which affects the intensity of the signals received by the robot eyes, is left constant. Each light lasts for a random amount of time between 100 and 200 seconds before disappearing and being replaced with a new source of the same kind, appearing in a random location.⁷

⁶As an additional detail, I did not allow a neuron to connect directly to itself — this seems to help reduce somewhat the possibility of pathological neuronal self-stimulation.

⁷The code as written allows also for multiple lights of each kind, and in future experiments it might be interesting to investigate the effect of a proportional majority of lights of a certain kind; but in this work I did not investigate this matter in depth.

2.3 Genetics

Each robot genome encodes, as a sequence of floating points between 0 and 1, the weights and the biases of its neural network (which are then rescaled to the respective range in order to obtain the corresponding phenotype). Whenever a robot dies (of hunger, of overheating, or of old age), a new one springs into being in a random position of the plane: in this way, the population size is kept fixed, and (at least in the non-bounded lifespan cases) the speed of evolution is made to depend on the overall level of fitness of the population.

If sexual reproduction is allowed, the new robot’s genome is obtained by selecting (through a roulette-type extraction) two of the longer-lived robots so far and performing a gene-by-gene crossover between them; otherwise, only one genome is extracted as the “parent” and selected. In either case, then the genome is then subjected to a round of mutation, during which each weight and bias has a 10% chance of being changed of a small amount.

If lateral gene transfer is allowed, each robot which is within ten energy points of dying of overheating or of hunger (but *not* of old age) and is younger than 200 seconds can undergo it with a fixed probability of 10% per second. If that happens to be the case, a “source” robot is selected through roulette-type extraction, each gene is overwritten with the source robot’s with 10% probability, and each gene is mutated with 10% probability. The overall effect of this mechanism is to accelerate the global evolution rate of the whole population whenever “young” robots are consistently failing to survive — or, in other words, whenever the circumstances of the environment have changed radically. In a break from biological plausibility, but in order to heighten the degree of similarity between this kind of gene transfer and Ashby-style ultrastability, I allowed the “new” genome of the robot to immediately overwrite the connection weights and the biases of its neural network.

It is worth remarking here that, despite the superficial similarity of their mechanisms, horizontal gene transfer and sexual reproduction operate quite differently. Sexual reproduction selects two old (that is, successful) genomes, creates a new one through crossover and mutation, and introduces a new robot in a random location with the initial amount of energy; gene transfer, however, operates on *unsuccessful* genomes by allowing them to “copy” genes from more successful ones and by increasing their rate of mutation; and as we will see in the next section, the consequences of the two mechanisms can be quite different indeed.

2.4 Protocol

I considered the three following main scenarios:

- I. Asexual reproduction, no lateral gene transfer;
- II. Asexual reproduction, lateral gene transfer;
- III. Sexual reproduction, no lateral gene transfer.⁸

For each scenario, I examined the four following possibilities:

- a) Unlimited lifespans, no extra noise when in danger;
- b) Limited lifespans, no extra noise when in danger;
- c) Unlimited lifespans, extra noise when in danger;
- d) Limited lifespan, extra noise when in danger.

For each of the resulting twelve possibilities, I used the following protocol:

⁸The “sexual reproduction and lateral gene transfer” possibility has not been examined in depth, partly because having both mechanisms of gene transfer seemed redundant and partly because of time constraints. At any rate, according to preliminary experiments there seemed to be little to no difference between this scenario and the “sexual reproduction, no lateral gene transfer” one.

1. Set the energy of both lights to 5000 (a value high enough for the robots to survive if they get close enough, or kill them if they get too close) and run the simulation for 30000 seconds.
2. If the robots have evolve to chase (at least) one light, set the energy level for *that* light to 500 and that for the other one to 10000. If both lights are being chased, decrease the energy emissions of the one which seems to attract the most robots and increase those of the other one. If instead no active searching behaviour is observable, run the simulation for 30000 more seconds without changing the parameters; and if no active searching behaviour has evolved by then, conclude that the current settings are not conducive to the development of simple homeostatic behaviour.
3. After 30000 seconds with the new energy setting, take note of the behaviour of the robots and swap the energy levels of the two light sources. Then run the simulation again for 30000 seconds.

In order to evaluate the behaviour of the robots, after each step I ran the simulation for 400 seconds, recording the activity of the robots for visual inspection (the resulting movies — recorded at an accelerated framerate, for ease of viewing — are saved in the supplementary material of this work) and then I recorded the distribution of their ages and energy levels.

For the configurations for which no seeking behaviour was observed even after 60000 seconds I repeated the whole protocol once more, for safety’s sake, with the intention of focusing on this second case if a seeking behaviour emerged. However, this never happened — and although the limited number of runs and population sizes have to be kept into account, this lends a some support to the reliability of these observation.

3 Results

In the following subsections I will summarize the results of the simulations for the various settings. I will do so first by presenting the results in a table, in which — for our different choices concerning noise and lifespan, and for each stage of the simulation — I will register if I observed the robots chase the red and blue lights strongly (**R**, **B**) or weakly (**r**, **b**). For any row, in the column corresponding to the first stage in which robots have been observed chasing at least one light, I underlined the light which they seemed to pursue the most; then in the next stage, I reduced the energy emissions for that light and increased those for the other one, while in the next one again I did the opposite.

These evaluations are inherently subjective, and I can only invite the reader to observe the attached movies by themselves. For each main scenario (but not for every setting) I will also present some more quantitative data about the distribution of robot ages and energy levels with respect to the subscenario which, in my opinion, presented the greatest amount of homeostatic behaviour.

I: Asexual reproduction, no gene transfer

This is the “baseline” scenario, whose results formed the basis for evaluating and comparing the effects of lateral gene transfer and sexual reproduction. The following table contains a summary of the main behaviours observed for the different regimes of this scenario:

Noise	Lifespan	Phase 1	Phase 2	Phase 3	Movie files
Constant	Unlimited	<u>r</u>	R	R , b	exp1part1-3.mp4
Constant	Limited	<u>R</u> , B	none	b	exp2part1-3.mp4
Varies	Unlimited	<u>r</u> , b	none	R , b	exp3part1-3.mp4
Varies	Limited	none	<u>R</u>	B	exp4part1-3.mp4

Scenario one: No sex, no lateral gene transfer

As we can see, even in this simple setting the robots seem quite capable of evolving to develop light-seeking behaviour; and as it can be seen in the accompanying movies, they also learn how to hover at a “safe” distance from the lights. However, it would seem that the robots have some difficulty adapting to environmental changes (in particular, switching focus from one source to the other in response to them).

Let us consider for example the first experiment: after Phase one, in which the robots succeeded — to some degree — in seeking the red light and staying at a safe distance from it, I reduced the energy released by it and increased that released by the blue light. But instead of switching to exploiting this richer source of energy, the robots simply decided to focus even more on the red source; and in Phase three, in which this source suddenly started releasing much more heat (and thus became far more dangerous), they kept mostly hovering around it, only increasing their distance from it.

The simulations of experiments two and three fared worse than this, as the robots substantially lost their light-seeking strategies altogether in response to the environmental changes of Phase two.

The last experiment of this first scenario was perhaps the most interesting one. The robots were not able to learn an adequate light-searching behaviour within their first 30000 seconds; but after 30000 more (with no energy intensity change), they developed the ability to seek the red light source. After I decreased the energy released from it and increased the one released from the blue source, they mostly switched to seeking preferentially the blue source; however, their overall behaviour (see the movie `exp4part3.mp4`) seemed to depend more on the overall intensity of the energy that they were receiving than by the precise color of it.

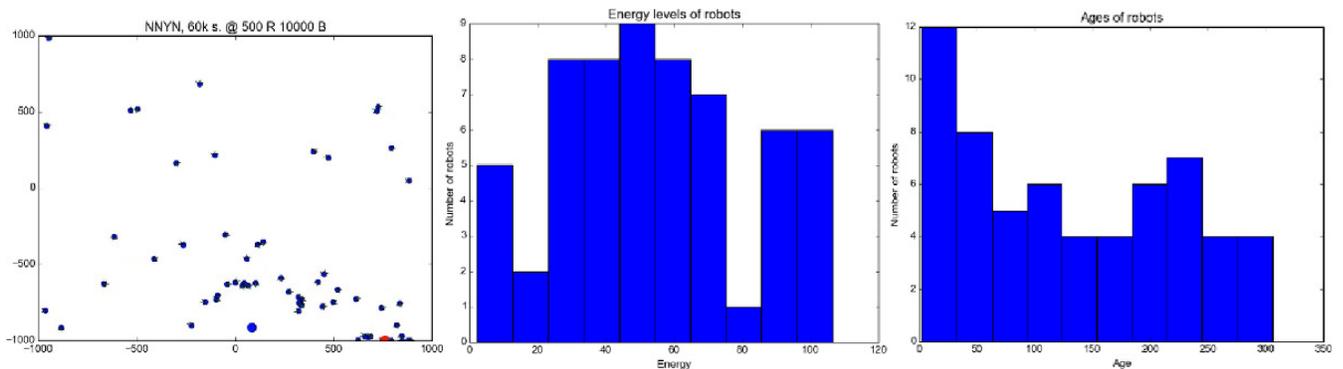


Figure 4: Experiment three, part three: energy and age distribution

In Figure 4, I represent the energy and age distributions of the robots at the end of this simulation. The energy distribution shows evidence of some degree of homeostatic behaviour, as the energy reserves of most robots lie close to the center of the viable interval; but this behaviour is far from general, and in particular we can see that a respectable number of robots are dangerously “hot” and in serious danger of overheating. This is possibly a consequence of the presence of danger-related noise: since robots with more than 1000 units of energy experience a sudden increase of perceptual noise, it is a viable strategy to wait for this “signal” and move away from the source only when they receive it; but further study would be required to confirm whether this is indeed the case.

The age distribution shows that several robots can be quite long-lived indeed, which confirms that this strategy is fairly reliable.

II: Asexual reproduction, lateral gene transfer

Noise	Lifespan	Phase 1	Phase 2	Phase 3	Movie files
Constant	Unlimited	<u>R</u>	B	r, b	exp5part1-3.mp4
Constant	Limited	<u>b</u>	R	b	exp6part1-3.mp4
Varies	Unlimited	<u>R, B</u>	R, B	R, b	exp7part1-3.mp4
Varies	Limited	<u>R,B</u>	B	R	exp8part8-3.mp4

Scenario two: No sex, lateral gene transfer

In this scenario we investigate the effect of lateral gene transfer. As we can see, this mechanism consistently increases the robot’s ability to achieve homeostasis and to adapt to changing circumstances: the robots

consistently rely on the most energetic light source (when not on both of them), and the energy distribution shown in Figure 5 is evidence of a well-developed homeostatic behaviour. Compared to the age distribution of Figure 4, the one of Figure 5 is more skewed towards younger ages; I suggest that one of the reasons for this is that lateral gene transfer is an inherently risky strategy - it can give an unsuccessful young robots the right tools for turning its luck around, but it can also damage its neural network beyond survivability (when doing nothing and relying on the old strategy a little longer could have been more successful): in other words, lateral gene transfer appears to trade short-term survivability for long-term adaptability.

It might be interesting to see how changing the cutoff age for lateral gene transfer, or the probability of post-gene transfer mutations, affects the overall evolution of the robots' behaviour; but this will be left to future works.⁹

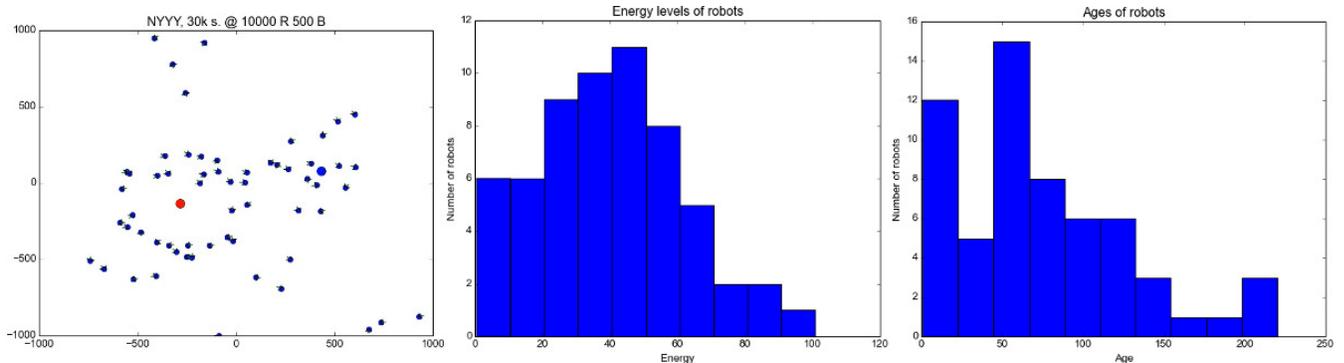


Figure 5: Experiment eight, part three

III: Sexual reproduction, lateral gene transfer

Noise	Lifespan	Phase 1	Phase 2	Phase 3	Movie files
Constant	Unlimited	none	none	X	exp9part1-3.mp4
Constant	Unlimited	none	none	X	exp10part1-3.mp4
Constant	Unlimited	none	none	X	exp11part1-3.mp4
Constant	Unlimited	none	r, b(both doubtful)	none	exp12part1-3.mp4

Scenario three: Sex, no lateral gene transfer

In this last scenario, I considered sexual reproduction as a tool for the development of homeostatic behaviour and the adaptation to changing circumstances. As we can see, the result was a clear failure: despite repeated attempts, the robots failed to develop a suitable light-seeking behaviour, opting instead for an essentially stationary strategy based on not moving much and hoping that a light would eventually appear at a suitable distance from them. In only one case weak light-seeking behaviours were observed; and, as it can be observed by watching the file `exp12part2`, it is a bit of a boundary case — robots mostly turn in place or wander randomly around, with not much in the way of purposeful activity. The final distribution of energies shows only a very weak level of homeostasis, with the energy levels almost uniformly distributed in the 0 – 80 range, and the age distribution is very heavily skewed towards the left, demonstrating an extremely high mortality rate.

Interestingly, if we remove the possibility of overheating, sexual selection allows the robots to develop very rapidly (within 1000 seconds) the ability for light-seeking behaviour, as shown in file `expExtra.mp3`. Thus,

⁹Additionally, in hindsight it would probably have been a better idea to record not the age distribution at the end of the simulation, but rather the overall lifespan length distribution during the last 1000 seconds or so of the simulation; and to plot the lifespan versus the mean energy levels of the corresponding individual through its lifetime.

it would appear that the fundamental weakness of sexual selection in this setting consists in its inability to preserve homeostatic behaviour across generations: if two robots employ different neural circuits to hover at a “safe” distance from a source, it is not particularly likely that their offspring will share this capability — for instance, it might contain *both* circuits, and hence tend to get far too close to the energy source.

It is possible, of course, that a different choice of crossover operator or a different genetic representation of the neural network could improve matters — this is quite likely, indeed, since a considerable number of animal species have relied on sexual reproduction to evolve impressively elaborated forms of homeostasis. But nonetheless, this result confirms that, under *some* circumstances, lateral gene transfer may be a far more suitable tool than sexual reproduction for the development of homeostasis and its preservation in the event of environmental changes.

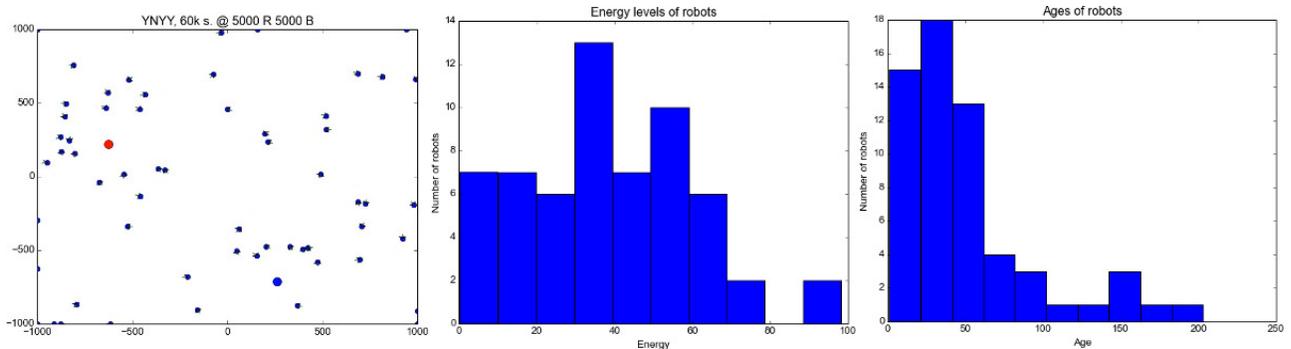


Figure 6: Experiment twelve, part two

4 Conclusions and future work

The main contribution of this work, as just mentioned, is a practical demonstration that (at least in *some* settings, such as the one of these simulations) bdelloid-style lateral gene transfer may prove very useful for the development of homeostasis and the adaptation to environmental changes, while sexual reproduction may be actively detrimental to it. This is merely a proof of concept, of course, and more research would be required in order to ascertain the *exact* circumstances that make asexuality plus this style of lateral gene transfer so evidently advantageous in the case of bdelloid rotifers.

Apart from this, more extensive (and computationally expensive) simulations and data collections would be necessary in order to infer more from this model. Furthermore, the model itself could be extended in various ways: for instance, it might be interesting to see the effects of greatly extending the size of the arena, allow for more lights (emitting also *mixtures* of red and blue lights, and make it so that lateral gene transfer can only copy genes from *nearby* robots.

In any case, the code developed for this project proved itself a suitable platform for the investigation of the evolution of homeostasis and adaptation, and the phenomenon of bdelloid-like lateral gene transfer (understood as a variant of Ashbian ultrastability) proved itself a suitable avenue for improving the degree of adaptability of a population to changing circumstances.

References

- [1] William Ross Ashby. *Design for a Brain.* , 1954.
- [2] William Ross Ashby et al. An introduction to cybernetics. *An introduction to cybernetics.*, 1956.

- [3] Chiara Boschetti, Natalia Pouchkina-Stantcheva, Pia Hoffmann, and Alan Tunnacliffe. Foreign genes and novel hydrophilic protein genes participate in the desiccation response of the bdelloid rotifer *adineta ricciae*. *The Journal of experimental biology*, 214(1):59–68, 2011.
- [4] Ezequiel A Di Paolo. Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions. *From animals to animats*, 6:440–449, 2000.
- [5] Diego Fontaneto and Timothy G Barraclough. Do species exist in asexuals? theory and evidence from bdelloid rotifers. *Integrative and comparative biology*, page icv024, 2015.
- [6] Eugene A Gladyshev, Matthew Meselson, and Irina R Arkhipova. Massive horizontal gene transfer in bdelloid rotifers. *science*, 320(5880):1210–1213, 2008.
- [7] Roberto Guidetti and K Ingemar JoÈnsson. Long-term anhydrobiotic survival in semi-terrestrial micrometazoans. *Journal of Zoology*, 257(2):181–187, 2002.
- [8] Olivia P Judson and Benjamin B Normark. Ancient asexual scandals. *Trends in Ecology & Evolution*, 11(2):41–46, 1996.
- [9] John Maynard Smith. Evolution: contemplating life without sex. *Nature*, 324:300–301, 1986.
- [10] Claudia Ricci. Life histories of some species of rotifera bdelloidea. *Hydrobiologia*, 104(1):175–180, 1983.
- [11] Claudia Ricci and Giulio Melone. Key to the identification of the genera of bdelloid rotifers. *Hydrobiologia*, 418(1):73–80, 2000.
- [12] BM Waggoner and GO Poinar Jr. Fossil habrotrochid rotifers in dominican amber. *Experientia*, 49(4):354–357, 1993.
- [13] David B Mark Welch and Matthew Meselson. Evidence for the evolution of bdelloid rotifers without sexual reproduction or genetic exchange. *Science*, 288(5469):1211–1215, 2000.
- [14] Michael James Denham White. *Modes of speciation*. San Francisco: WH Freeman 455p.-Illus., maps, chrom. nos.. General (KR, 197800185), 1978.