

Evolving Collective Behaviours in Simulated Kilobots

Jane Holland, Josephine Griffith, Colm O’Riordan
National University of Ireland, Galway
{j.holland3, josephine.griffith, colm.oriordan}@nuigalway.ie

Abstract

The field of Evolutionary Robotics has multiple common tasks and widely used benchmark activities such as navigation, obstacle avoidance, and phototaxis. We present an evolutionary approach to learning behaviours that demonstrate emergent collective phototaxis in a swarm of simulated robots.

Our approach demonstrates that evolutionary computation can be used to evolve the emergent, self-organising behaviours of clustering and phototaxis in a population of simulated robots where the robots possess limited capabilities. In addition to demonstrating the feasibility of the approach, we show that the evolved behaviours are also robust to noise and flexible in changing environments.

1 Introduction

Evolutionary swarm robotics is a field that uses evolutionary computation techniques to evolve behaviours for robots in multi-robot systems. These multi-robot systems emphasise decentralised and self-organising behaviours. These behaviours are an indirect result of local sensing and local communication amongst individuals with limited abilities. Typically, these multi-robot systems, whose behaviours are inspired by swarms in nature – ants, bees and termites – are physically quite simple. Similar to insect societies, robotic systems need to be robust, scalable and flexible to changing environments. These desired features can be achieved through different forms of self-organisation such as aggregation, coordination and exploration. Self-organisation is a process whereby a system transitions from an unordered state to an ordered state through communication, be it between agents or between agents and their environment.

Collective phototaxis is the ability of a group of agents to sense a light source and move, as a collective, towards (or away from) the light source. In order to achieve this coordinated behaviour efficiently, agents must be able to communicate amongst each other in the group; this can be accomplished by either direct and indirect communication. Direct communication refers to the exchange of information without the need for physical interaction, whereas indirect communication refers to communication achieved by actions taken by agents in the environment. Hence, the environment

used in an experiment is important as both the intrinsic (swarm size and structure) and extrinsic (time, predation) factors used can heavily influence the overall results.

In this paper we describe the evolution of the behaviours of clustering and phototaxis in a group of simulated Kilobots (Rubenstein et al., 2012). The main contribution of this work is the successful demonstration of the evolution of collective phototactic behaviours in simulated robots that have limited computational and memory resources. Furthermore, we explore the robustness of these evolved behaviours by evaluating different parameter settings, by adding noise to the robots' light and communication sensors and by altering the experimental environment. The outline of this paper is as follows: initially, we briefly discuss related work in section 2. We describe the adopted methodology in section 3 which gives an overview of the robots whose simulated versions are used in this work (3.1) as well as describing the evolutionary algorithm, fitness evaluations (3.2) and experimental setup (3.3). In section 4 we discuss the results of our simulations: positive and negative phototaxis, and phototaxis in changing environments. Lastly, conclusions and future work are discussed in section 5.

2 Related Work

Much of the previous work in this area deals with the evolution of a control architecture with a high representational power to carry out simple tasks. For instance, Watson et al. (Watson et al., 2002) combined embodied evolution and neural networks to develop robot controllers that exhibited phototaxis from any location in an arena. Although this work was successful, it did not have any explicit robot interaction. In a similar vein, Francesca et al. (Francesca et al., 2014) also used neural networks. Robots were evolved to go to, or from, a light source if perceived; otherwise they would continue in a straight line. Additionally, the robots used a range and bearing system to detect the presence of neighbouring robots. Although Silva et al. (Silva et al., 2015; Silva et al., 2016) also evolve artificial neural network controllers in order to carry out phototaxis, they do so through online learning. These experiments use a distributed and decentralised neuroevolution algorithm that continuously modifies robots' behaviours in order to respond to changes and unforeseen circumstances.

Evolving spiking neural networks for phototaxis is a commonly adopted approach (Di Paolo, 2002; Batllori et al., 2011; Jeanson and White, 2012). A light source was placed in the centre of an arena and a robot's fitness was based on the average distance to the centre of the light source over its lifetime. According to Di Paolo et al. (Di Paolo, 2002), robots that were evolved using spiking neural networks demonstrate a rich variety of behavioural strategies in comparison to robots that were evolved using a neural network.

Roli et al. (Roli et al., 2015) explored positive and negative phototaxis through the use of Boolean networks. Their experiment used E-puck robots to demonstrate positive phototaxis until a sharp sound was perceived, after which they switched to demonstrate negative phototaxis. In this way, they study how robots adapt to new operational requirements without changing the experimental environment.

Although it has been previously shown that phototaxis can be evolved using neural

networks in robots with range and bearing systems, it has not been demonstrated using a different architecture with less representational power, nor using robots with such a degree of constraints.

3 Methodology

Although it may seem that exploring the evolutionary process of phototaxis has already been well explored (Nelson et al., 2009), our work differs from previous approaches in a number of ways. In contrast to the work discussed in the related work section, we evolve collective phototactic behaviours using local interactions between robots that do not have any bearing systems. Moreover, the robots have *no a priori* knowledge at the beginning of the experiment; robots are unaware of the location of their neighbours or of the light source. This makes the learning task more difficult.

Through the use of infra-red distance and ambient light sensors, the robots can measure their distance from neighbouring robots as well as from the light source. The robots' evolved behaviours take these measures as inputs and use them to set the values for the left and right motors. The evolutionary process learns the settings for these motors for the corresponding distance measures.

We diverge from the traditional neural network control architecture and instead use a look-up table to represent the robots' behaviours in various settings. The values in this table are subject to evolutionary pressures. By adopting a lookup table, the memory constraints of the robots do not become an issue. Decentralised algorithms do not require a leader and so, are often a more robust and flexible learning method over others. In other words, our model does not require robots to complete tasks in response to orders or requests from other robots. The group behaviours witnessed are a result of the individual behaviours. However, we do adopt a centralised learning approach using a single population genetic algorithm. In this method useful genotypes are propagated, and are subject to mutation and crossover. This approach contrasts with explicit decentralised learning mechanisms where individual robots learn from their own individual experiences.

Using the approach of learning values (stored in a lookup table) for various environmental settings, we show the evolved behaviour of phototaxis. Following on from this, we test the robustness of the behaviours by changing the learning environment throughout the task and adding noise to the robots' light and communication sensors.

3.1 Kilobot

The Kilobot, shown in Figure 1, is a low-cost robot designed to make testing on collective behaviours scalable and easily comprehensible. This robot uses an Atmega328 single-chip microcontroller (8MHz with 32K of memory) to run user-defined programs as well as operate as an interface between motors, power circuitry and LEDs. The Kilobot is not a conventional robot as it uses vibration motors for locomotion and reflective infra-red light to communicate with neighbouring robots. Two pulse width modulation channels are used to regulate the speed of the motors, which result in noisy feedback and thus, prevent the Kilobot from travelling long distances. However, by using




Figure 1: Bottom (left) and isometric (right) views of the Kilobot. Key features include: vibration motors, infra-red transmitter/receiver, LED, and charging tab.

measured distances with distance sensing as feedback, the robot’s movement can be realigned and be used in promoting the use of collective behaviours (Rubenstein et al., 2012).

Coordinated navigation can be difficult to implement in these robots as they lack a bearing system and have a limited communication range of up to 10cm. Furthermore, they can only perceive the distance between each other, but not the direction in which their neighbours are travelling. To overcome this complication, Kilobots continuously measure the intensity of the incoming infra-red signals from their neighbours through 10-bit analog-to-digital converters. The infra-red transmitter and receiver allow the robot to receive messages from every direction within range. This transmitter works by emitting light to a neighbouring robot by means of reflecting the signal off the surface on which the Kilobot has been placed.

The simplicity of this robot can be used in our favour as basic forms of communication are enough to accomplish collective behaviours, and can also be easily scaled up with the number of individuals involved. Furthermore, although evolutionary computation has been performed on Khepera robots (Mondada et al., 1994) and S-Bots (Mondada et al., 2002), there has been no research carried out on robots with a high level of hardware and software restraints.

In order to carry out our experiments, we used a realistic model called ARGoS; an open-source 3D world, multi-robot simulator (Pinciroli et al., 2011) (Figure 2). This tool has Kilobot hardware compatibility; it provides Kilobot models with all the functionalities available on real Kilobots. Furthermore, this simulator accurately models physical laws such as mass, friction and acceleration. Thus, ARGoS can accurately simulate different sensor devices such as infra-red proximity and LED sensors.

3.2 Evolutionary Algorithm

Evolutionary Algorithms are very useful for synthesising efficient self-organizing behaviours that are simple, robust and flexible to changing environments (Trianni and Nolfi, 2011). The features of evolutionary algorithms can make problem solving somewhat easier as evolutionary algorithms can decrease the computational cost as well as




Figure 2: A graphic visualisation of the experimental arena designed with ARGoS.

the time required to solve a problem. Furthermore, evolutionary algorithms have been successful in finding novel behaviours that are often counter-intuitive. It is for these reasons that we chose to use evolutionary algorithms in our approach. We rely on the evolutionary process to tune the dynamics of the interaction between the robot and its environment. Generally, a control architecture is used that has a high representational power, such as a neural network. However, this type of architecture can prove counter-productive in uncertain and dynamic conditions where there are multiple robot to robot interactions (Francesca et al., 2014). Thus, as a consequence, the agents may exploit idiosyncratic features of the environment to increase their fitness, which leads to a decrease in performance in new environments where these features are absent (Floreano and Keller, 2010).

We use a look-up table (LUT) to evolve 10 left and right motor speeds. These values are evolved using a genetic algorithm; these values determine how our controller reacts in the environment. The controller of the robot can be seen as a phenotype of the robot and it determines the robot’s behaviour; it is this behaviour which is assigned a fitness value. The phenotypes are represented by genotypes that undergo evolutionary operators. Each Kilobot holds a vector of different actions which are linked to the robot’s actuators i.e. speeds between 0-1 cm/s. When a Kilobot perceives a neighbouring robot, or senses a light source, it looks-up a specific action in the LUT (Table 1) in order to respond to the current situation. Each row in the lookup table specifies the values for the left and right motors for particular environmental situations. Each robot potentially evolves a different LUT due to the presence of noise and mutation.

The initial population is composed of 50 random genotypes, which are encoded into, and mapped to, the robot’s controller. Each run of the experiment lasts 150 or 200 generations, depending on the task carried out. The genetic algorithm adopted is a simple generational model with tournament selection, elitism, crossover and mutation. Once every genotype of the population has been evaluated by means of a fitness function, two individuals are picked at random from the population and the fittest of the pair is chosen to act as a parent. Chosen parents produce offspring until the subsequent generation is full. The offspring are formed through a process of single-point crossover; the crossover rate is 30% and are subject to a 2% mutation. The process is repeated until suitable behaviours is evolved.

Kilobots are rewarded a higher fitness when the distances between neighbouring

Table 1: An example of a LUT with learned Left and Right motor speeds.

Left Motor	Right Motor
0.8958304382	0.8643559063
0.4791235448	0.6309483493
0.2605420978	0.4704560627
0.0612328765	0.8783525228
0.1255781651	0.8786792431
0.5989533024	0.8958941174
0.2688139217	0.1218344814
0.9344603782	0.7710112717
0.9582202760	0.8889869242
0.8249491609	0.0645951105

Table 2: Parameter settings used in experiments

Acronym	Definition
O	Original <i>posP</i> (equation 2) or <i>negP</i> (equation 3) used depending on the experiment.
N	Light reading values are normalised to be in the same range as the distance readings.
D	If more than one robot is perceived, then the average distance of all robots in this proximity is taken.
M	Offspring are subject to 2% mutation (the probability of changing a motor value after crossover takes place).

Different combinations of these parameters are used throughout the experimentation. For instance, **ODM** refers to the use of the original phototaxis equation while implementing the average distance, and applying mutation.

robots are reduced and their light reading is decreased (or increased, depending on the experiment). Therefore, the more dense a group and the closer it is to (or further from) a light source, the higher the individual fitness, and in turn, the group fitness. For each evaluation, the approach remains the same; each robot reads in a *distance* of 1-100mm and a *light* reading between 0-900cd. Depending on the experiment, the value of either *posP* (towards the light) or *negP* (away from the light) is used as a parameter to get the *index* of our LUT; this is where the left and right motor values are stored. The returned index is then evaluated via the power law (*P*) and the robot's fitness (*Fe*) is accumulated:

$$Fe+ = P(\text{getIndex}(\text{posP})) \quad (1a)$$

$$P = \text{maxK} \cdot \text{index}^{-\alpha}, \quad (1b)$$

where *maxK* is the maximum local performance of each robot; a constant factor. Scal-

ing the *index* by *maxK* causes a proportionate scaling of *P*. In this way, robots are evolved to reduce their speed if their neighbours are in close proximity, or if the light reading is increased (or minimised). Therefore, the collective behaviours of clustering and phototaxis can emerge.

The positive (*posP*) and negative (*negP*) phototaxis equations used are as follows:

$$posP = \sum_{i=1}^n (distance_i) + light \quad (2)$$

$$negP = \sum_{i=1}^n (distance_i) + light^{-1}, \quad (3)$$

where *n* is the number of interacting robots.

Furthermore, we use different combinations of parameter settings in each task to examine the robustness of the approach; these are referred to as variations of *O*, *N*, *D* and *M* throughout the paper and are described in Table 2.

3.3 The Experimental Setup

This paper describes four experiments which are outlined in two sections. The first section describes two experiments, positive phototaxis (Exp 1a) and negative phototaxis (Exp 1b) with one light source in an unvarying environment. The second section describes positive phototaxis experiments (Exp 2a and Exp 2b) in environments that change throughout the duration of the task.

3.3.1 Positive and Negative Phototaxis

In both the positive (Exp 1a) and negative phototaxis (Exp 1b) experiments, robots are randomly placed (in terms of both position and orientation at the beginning of each generation) in a 1m×1m enclosed arena with a light source placed in the centre (Figure 3A). Each experiment has 10 trials each of which last 500 seconds, with 10 ticks per second (5000 seconds in total). The population size is constant and set to be 50 and the crossover type is always single-point. In each trial, the behaviours of the robots are evolved to either cluster under the light source as a collective group (positive phototaxis), or to move away from the light source while attempting to stay in a cluster (negative phototaxis). Uniform noise was added to the robots' light sensors in the range of [-1,1] with a probability of 20%; the final sensor reading is normalised in the [0-1] range. Gaussian noise was also added to the communication sensor (standard deviation $\sigma = 1$). Both of these experiments explored and compared the outcomes for different parameter settings in order to discover which resulted in the best evolved behaviours, for example, those that achieved the best clusters near the light or those that achieved these outcomes in the most efficient way; this is explained further in section 4. Parameters tested are: normalisation of light readings (N), average distance (D) and mutation (M).




Figure 3: Graphic visualisation of arenas used for experimentation. A: The light source is situated in the centre of the arena (Exps 1a and 1b). B: The light positioned in the top left hand corner is turned off during the task (Exp 2a). C: The light constantly moves at 1 micro metre every 10 seconds; it begins at the bottom right hand side and progresses towards the top left hand side (Exp 2b).

3.3.2 Phototaxis in a Changing Environment

As with the positive and negative phototaxis task described in the previous section, we use a $1\text{m} \times 1\text{m}$ enclosed arena, but with a varying light source. There were two variations considered in this paper: in the first version, we turned off one light source as shown in Figure 3B (Exp 2a) and in the second version, the light position constantly changed as shown in Figure 3C (Exp 2b). All changes occurred throughout the duration of the experiments. As the second variation is a more difficult landscape to learn, we increase the experiment length to 1000 seconds, but the ticks per second remain at 10 (10000 seconds in total) and the population size stayed at 50. These tasks required adaptive behaviour in order to successfully carry out positive phototaxis. Again, we explore different parameter settings in each task. Furthermore, uniform noise was added to the robots' light sensor in the range of $[-1,1]$, and normalised. The level of Gaussian noise to the communication sensor remained the same.

4 Results and Discussion

Experimental outcomes were evaluated using two separate measures or criteria: average group performance over 10 trials and graphic visualisations of the resulting clusters. Furthermore, we present the evolved motor speeds. The group performance is the summation of individual fitnesses in each generation which allows us to compare the fitness of the populations for different parameter combinations. The visualisations of the resulting clusters gives an overall indication of how well the robots cluster together and how well they exhibit phototaxis. Furthermore, we plot the values of the motor values at different generations throughout the evolutionary run in order to discover whether there is any correlation between group performance and motor speeds for both left and right motors.




Figure 4: Average group performance of all parameter settings in Exp 1a after 10 trials.

4.1 Positive Phototaxis

We examined six combinations of the parameters in this experiment: O , OM , ODM , N , NM , NDM . The experiment ran for 150 generations. As seen in Figure 4, it appears that O , OM and ODM are the best for group performance settings. However, their high scores are due to the non-normalisation of the light readings and are, in fact, the settings with the worst performing measures. We illustrate ODM in Figure 5 to demonstrate this argument.

Although six parameters settings were tested, we focus on the top two performing parameter combinations, N and NDM . The evolutionary process for the parameter setting, N (normalised distance to the light) appears to stabilise before the 20th generation and yet, gives a slightly lower average group performance than when both mutation and distance values (NDM), were considered. In contrast, NDM takes longer to maintain a steady group performance (achieved by generation 50); this may be due to the mutation rate.

When visually examining the simulations, it appears that the parameter setting N demonstrates the best overall collective phototactic behaviours (Figure 5) in comparison to NDM .

In terms of the evolved motor speeds, on average, the experiment with parameter setting N had a higher speed than NDM . However, through the course of the evolutionary runs, both motors continued to slow down; starting at 0.51 – left – and 0.49 – right – and dropping to 0.46 and 0.33 respectively (Figure 6). In contrast, the parameter settings NDM led to a speed slower than observed in the experiment with setting N , due to a dip by the 50th generation. This seems to suggest that the more a group is clustered on average, the lower the motor speeds.

Although it could be argued that phototaxis was evolved, dense clusters were not exhibited for the settings O , OM and ODM . Despite these outcomes, these settings resulted in the highest group performance of the parameter settings tested. It appears that the high group performance score may, in fact, be due to the addition of the distancing measure as well as the non normalisation of the light readings. The robots




Figure 5: Graphic visualisations evolved positive phototaxis after 150 generations (Exp 1a).




Figure 6: Evolved motor speeds of the top three parameter settings in Exp 1a. after 10 trials

appear to form a number of ‘loose’ clusters. This suggests that the distance measure is out-weighting the influence of the light factor as robots no longer cluster towards the centre. Therefore, this setting still receives a high group performance as the distance is reduced and the robots are still in a cluster, albeit, not in the expected formation.

When comparing the outcomes for the *O*, *OM* and *ODM* settings, they all result in a successful evolution of a collective positive phototaxis behaviour, despite being loosely clustered. The only variation between these parameters settings with *N* *NM* and *NDM* is the normalisation of the light readings.

With regard to motor speed, *ODM* had the highest overall speed, 0.11 times quicker than *N*. The left motor speed came to a plateau by generation 50, and the right motor by generation 100. This led to a more dominant right spin as the left motor speed has a larger decline in comparison. In contrast, all settings except *ODM* and *NDM* evolved a left spin.




Figure 7: Average group performance of all parameter settings in Exp 1b. after 10 trials

4.2 Negative Phototaxis

These experiments also ran for 150 generations and negative phototaxis was successfully evolved in all cases. The parameter settings tested remain the same as in Exp 1a. Again, we discuss the top performing parameter settings; *ODM* and *OM*.

Experiments with the parameter settings *ODM* and *OM* obtained the best overall group performance scores. For both of these settings, the evolutionary runs resulted in populations that appear to obtain close to an optimal level; this high performance is achieved very early on in the evolutionary run of *ODM*, which continues to increase after the 25th generation. Other parameter settings also exhibited a good performance (Figure 7).

The evolution begins to converge by the 45th generation for the parameter setting *ODM*; this was considerably faster than observed in the setting *OM* which steadily evolved slowly and began to stabilise by generation 65 (perhaps due to the absence of the distance measure in the fitness function). The distance measure in *ODM* in each generation results in the creation of outliers as witnessed in Figure 8. Despite these outliers, experiments with this parameter combination leads to behaviours that performs quite well with, on average, 84% of the robots, successfully carrying out negative phototaxis in a group. The existence of the outliers may also explain why using the parameter setting *ODM* resulted in a better group performance score than that in the experiments with setting *OM*. These outliers are still within close proximity to each other due to the distance measure and so, they lead to a better overall score.

Upon examining the evolved motor values, it is clear that mutation does not have a huge effect. Both settings (*O* and *OM*) resulted in similar outcomes, both of which still evolved 60% of the same motor directional trajectory and both parameter settings evolved to travel around the arena in a clockwise manner. However, the setting *OM* resulted in speeds that were 0.87 times slower on average that those obtained with setting *ODM* (Figure 9).




Figure 8: Graphic visualisations of evolved negative phototaxis after 150 generations (Exp 1b).




Figure 9: Evolved motor speeds of the top two parameter settings in Exp 1b.

4.3 Phototaxis with a Changing Light Source

In this experiment an additional light source was introduced; one light was placed at each of two diagonally opposite corners. Initially, both lights were turned on; they remain on for the first 2000 seconds of the task, after which one was turned off. In the first 2000 seconds, robots are rewarded for clustering near either light source. Following the change in environment, there is now only one light source to cluster around and robots with the correct evolved behaviours will move to the only light source remaining. This more complex environment is adopted to explore the robustness of the evolved behaviours. When we initially ran this experiment with the positive phototaxis equation (equation 2), it became apparent that too much weighting was given to the inter-agent distances in comparison to the light intensity factor. That is, although the distance between each robot and the distance from the light are taken into account, moving towards a light source will inadvertently bring robots closer to each other, thereby increasing the score related to the distance between neighbours. In an effort to correctly balance between the two factors, we chose to use the average rather than the




Figure 10: Average group performance of all parameter settings in Exp 2a. after 10 trials

summation of the inter-agent distances.

The experiment ran for 150 generations. We tested six parameter settings in this task: *O*, *OM*, *ODM*, *N*, *NM*, *NDM*. Similar to Exp 1a, it appears that *O*, *OM* and *ODM* are the best for group performance settings (Figure 10), but this is again due to the non-normalisation of the light readings leading to higher fitness scores. In reality, they have not performed as well as the the settings. In this section, we concentrate on the three best parameter settings: *NDM*, *NM* and *N*.

The *N* and *NM* settings demonstrate stabilisation early on during he evolutionary process in contrast to *NDM*, which continues to increase in performance up until around generation 30. Despite the continued fluctuation, *NDM* gave an average group score 0.21 and 0.17 times better than *N* and *NM* respectively.

When visually examining the simulations, we can see that *NDM* has two clear groups of clusters by the 2000th tick in comparison to *N* and *NM*. However, by 5000 ticks, the parameter setting *NM* demonstrates the best overall collective phototatic behaviours (Figure 11). Regardless, all three parameter settings lead to good performance with very few outliers.

With regard to evolved motor speeds, parameter setting *N* results in the fastest speeds (Figure 12). The only prominent difference in *NDM*'s motors speeds is its greater likelihood of a higher left spin. As this setting performed the best, it would suggest that an average speed of 0.49 with a marginal increase in power for the left motor is beneficial for the evolution of the correct group phototatic behaviour.

4.4 Phototaxis with a Changing Light Position

In this experiment, the light source is moved throughout the evolutionary run and we explore the power of the approach to evolve suitable behaviours when the environment is subject to change. This continuous movement experiment tested the same six pa-




Figure 11: Graphic visualisations of evolved positive phototaxis in Exp 2a. Images on the top line illustrate the evolution state at 2000 seconds, just as the light source on the top left is turned off. The images on the bottom represent the evolved behaviour at 5000 seconds.

rameters as in Exp 2a. Throughout the task, a light source is moved diagonally upward through the arena. This experiment lasted 200 generations in order to allow sufficient time for evolution. We again utilise the average of the inter-agent distances in the fitness function.

We focus on the best and worst overall performing parameter settings: *NDM* and *N* as both exhibit good collective phototactic behaviours. *NDM* exhibited the best overall collective phototaxis in terms of graphic visualisation (Figure 14) as it successfully clustered under the light source throughout the task and group performance (Figure 13) due to its high score.

N converged very early on during the evolutionary run, at 15 generations in comparison to *NDM* which stabilised at the 40th generation. Furthermore, although *NDM*'s setting had a group performance 0.3 times larger than *N*, both settings appear to have a evolved good clustering behaviours, which is illustrated in Figure 14. This would suggest that our approach is robust to changing environments and different types of noise.

When looking at the evolved motor speeds for this experiment (Figure 15), *NDM* takes a sharp decline in the evolved left motor speed in comparison to the right motor, which have a steady value of 0.49. Interestingly, the motors evolved from the *N* setting have a similar plateau, but with the left motor and a less sharp decline with the right motor. This would suggest the the ideal directional trajectory is to travel in an clockwise direction using an average power setting.

In summary, we successfully evolved phototaxis in our experimental trials. Generally, parameters that used the normalisation of light measure (*N*) performed better and used a slower motor speed in the positive phototaxis experiments. Contrary to this, in our negative phototaxis experiment, the original settings (*O*) had a much greater collective phototactic ability. Overall, our approach can be deemed as robust as the evolved behaviours have the capability to adapt to different environments, even with the addi-




Figure 12: Evolved motor speeds of the top three parameter settings in Exp 2a.




Figure 13: Average group performance of all settings in Exp 2b. after 10 trials




Figure 14: Graphic time-lapse visualisation of evolved positive phototaxis every 2500 seconds in Exp 2b.




Figure 15: Evolved motor speeds of N and NDM parameter settings in Exp 2b.

tion of mutation and different types of noise.

5 Conclusions

Experimentation has demonstrated that using evolutionary computation on simulated robots possessing limited capabilities is feasible. We have successfully evolved a swarm of simulated Kilobots to carry out collective positive and negative phototaxis as well as light following in a changing environment with added noise to the light and communication sensors.

In summary, we successfully evolved phototaxis in all of our experiments. Generally, parameters that used the normalisation of the light value performed better and used a slower motor speed in the positive phototaxis experiments. In contrast, the original setting had a much greater collective phototactic ability in our negative phototaxis experiment.

In future work, we wish to delve further into the interplay of the light and distance measures as it demonstrates interesting effects on the landscape. Lastly, we would like to carry out our experiments on physical Kilobots and discover other emergent self-organising behaviours.

6 Acknowledgements

The authors acknowledge the support of Ireland’s Higher Education Authority through the IT Investment Fund and ComputerDISC in the National University of Ireland, Galway.

References

- Batllori, R., Laramée, C. B., Land, W., and Schaffer, J. D. (2011). Evolving spiking neural networks for robot control. *Procedia Computer Science*, 6:329–334.
- Di Paolo, E. (2002). Spike-timing dependent plasticity for evolved robots. *Adaptive Behavior*, 10(3-4):243–263.
- Floreano, D. and Keller, L. (2010). Evolution of adaptive behaviour in robots by means of darwinian selection. *PLoS Biol*, 8(1):e1000292.
- Francesca, G., Brambilla, M., Brutschy, A., Trianni, V., and Birattari, M. (2014). Automode: A novel approach to the automatic design of control software for robot swarms. *Swarm Intelligence*, 8(2):89–112.
- Jeanson, F. and White, A. (2012). Evolving axonal delay neural networks for robot control. In *Proceedings of the 14th annual conference on Genetic and evolutionary computation*, pages 121–128. ACM.
- Mondada, F., Franzi, E., and Jenne, P. (1994). Mobile robot miniaturisation: A tool for investigation in control algorithms. In *Experimental robotics III*, pages 501–513. Springer.
- Mondada, F., Pettinaro, G. C., Kwee, I., Guignard, A., Floreano, D., Gambardella, L., Nolfi, S., Deneubourg, J.-L., and Dorigo, M. (2002). Swarm-bot: A swarm of autonomous mobile robots with self-assembling capabilities. In *Proceedings of the Workshop on Self-Organisation and Evolution of Social Behaviour*, number LIS-CONF-2002-007, pages 11–22.
- Nelson, A. L., Barlow, G. J., and Doitsidis, L. (2009). Fitness functions in evolutionary robotics: A survey and analysis. *Robotics and Autonomous Systems*, 57(4):345–370.
- Pinciroli, C., Trianni, V., O’Grady, R., Pini, G., Brutschy, A., Brambilla, M., Mathews, N., Ferrante, E., Di Caro, G., Ducatelle, F., et al. (2011). Argos: a modular, multi-engine simulator for heterogeneous swarm robotics. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pages 5027–5034. IEEE.
- Roli, A., Villani, M., Serra, R., Benedettini, S., Pinciroli, C., and Birattari, M. (2015). Dynamical properties of artificially evolved boolean network robots. In *Congress of the Italian Association for Artificial Intelligence*, pages 45–57. Springer.
- Rubenstein, M., Ahler, C., and Nagpal, R. (2012). Kilobot: A low cost scalable robot system for collective behaviors. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 3293–3298. IEEE.
- Silva, F., Correia, L., and Christensen, A. L. (2016). Leveraging online racing and population cloning in evolutionary multirobot systems. In *European Conference on the Applications of Evolutionary Computation*, pages 165–180. Springer.

- Silva, F., Urbano, P., Correia, L., and Christensen, A. L. (2015). odneat: An algorithm for decentralised online evolution of robotic controllers. *Evolutionary Computation*, 23(3):421–449.
- Trianni, V. and Nolfi, S. (2011). Engineering the evolution of self-organizing behaviors in swarm robotics: A case study. *Artificial life*, 17(3):183–202.
- Watson, R. A., Ficici, S. G., and Pollack, J. B. (2002). Embodied evolution: Distributing an evolutionary algorithm in a population of robots. *Robotics and Autonomous Systems*, 39(1):1–18.