

Animal-guided evolutionary computation in honeybees and robots

Pedro Mariano¹, Ziad Salem², Rob Mills¹, Payam Zahadat², Luís Correia¹ and Thomas Schmickl²

¹BioISI, Faculdade de Ciências, Universidade de Lisboa, 1749-016 Lisboa, Portugal

²Artificial Life Laboratory, Department of Zoology Karl-Franzens-University, Universitätsplatz 2, Graz, 8010, Austria
plmariano@ciencias.ulisboa.pt

Abstract

In this paper we report our ongoing work with evolving bio-hybrid societies. We develop robots that will be integrated in an animal society and will be accepted as a conspecific. Moreover, we want our robots to affect the behaviour of animals. We are using evolutionary algorithms to optimise robot controllers, where fitness is evaluated via measuring the effect a robot controller has on the animals. Several issues have to be considered: if the animals do not have a homogeneous behaviour several evaluations are needed to rule out outliers, and yet evaluating animal behaviour is a time consuming task. Besides the time it takes to record their behaviour, we have to take into account animal resting time, stimulus habituation, and feeding periods. Another factor that increases the task difficulty is robot heterogeneity, which is similar to the so called reality gap problem that occurs in evolving robot controllers in simulation. In our case, if we want a robust robot controller, we have to evaluate it in different robots. Overall, we found that doing online on-board evolutionary computation with robotic devices and animals is extremely challenging and we provide clues to avoid its major pitfalls.

Introduction

Studying animal interaction by exploiting artefacts has a long and rich history in biological research (Lorenz, 1950; Tinbergen, 1948). In past decades dummy objects were presented to animals either statically or following a fixed movement pattern to trigger behavioural responses from these animals. The basic idea is that those artefacts emit key signals, mostly visual stimuli, to animals in order to study the signals that they use for interacting with their species mates. In social insects, where swarm intelligence arises from simple social interaction patterns in an emergent way, the study of those key patterns of interaction is highly significant, for example to study the coordination in honeybees (Seeley, 1995).

In recent years the progress in the engineering and miniaturization of computational devices and mechatronic artefacts allowed those studies to leap to a new level of complexity, as today's dummy devices can exhibit a life of their own. They can sense the animals through their sensors, decide on their next actions based on a behavioural model and then

perform behaviours through their actuators. This way, for the first time in history, dummies can enter into a dialogue with the tested animal. However, often the animal is not even tested: in contrast, it is the behavioural model implemented onto the robotic dummies that is tested for acceptance by the animals. Thus, in this novel line of research, the animal becomes a referee. This research has yielded multiple robotic implementations recently, ranging from honeybee-dancing robots (Landgraf et al., 2011) to robots that integrate themselves into fish shoals (Bonnet et al., 2016; Landgraf et al., 2016; Donati et al., 2016) and to robots that interact with cockroaches (Caprari et al., 2005).

However, all these approaches resorted to the idea that an *a priori*-available behavioural model is given to the robotic devices. In the approach we pursue with our ASSISIBf project,¹ we aim to get rid of this *a priori* model and use computational methods to develop the model during runtime, following ideas that were recently suggested for robotics (Cully et al., 2015). As a first step, we here provide the robots with a key signal known from honeybees, the “queen piping signal” (Michelsen et al., 1986a). We abstract this signal to a simple series of frequency pulses intermittent by a significant break. We identify two key parameters that determine the effect of the signal, frequency and pause length, and use evolutionary computation to “evolve” a signal that has a maximum effect on young honeybees’ motion behaviours. In order to provide fitness feedback, we observe the honeybees by automated video analysis and feed back the behaviours into the evolutionary loop as a fitness value.

Although this system is very simple, it provided us with a rich set of pitfalls and caveats to be learned, into which we happily stepped and fell during the course of our experimental procedure. This article’s aim is to explain those pitfalls to allow other scientists to avoid them, as they are stepping into this novel emerging field of bio-hybrid robotic research. Ultimately we identify the best approaches we found so far, in order to provide a first recipe for this novel type of research.

¹<http://assisi-project.eu/> See also Schmickl et al. (2013a,b) for an overview.

Related research

To the best of our knowledge this is the first work on evolutionary algorithms (EAs) with fitness provided by animals other than humans. EA with fitness provided by humans is an idea first proposed by Dawkins (1986) and later used in a variety of applications, from graphic art to engineering, under the common designation of *Interactive Evolutionary Computation* (see Takagi, 2001 for a comprehensive review).

The concept of so-called “mixed societies”², involving living organisms and artificial devices, can be traced back to the work of Goumopoulos et al. (2004) integrating plants with sensors and actuating gadgets, and shortly after integrating cockroaches and robots (Caprari et al., 2005). However, no evolutionary experiment had been attempted so far in such societies. It is interesting to note that ultimately we may consider evolution applied in both populations, which is tantamount to co-evolution (see Funes et al., 1998 for an experiment in a simulated setup where the animals are humans). In our evolving bio-hybrid society the animals do not evolve, since they are virtually of a single generation, while the robot controllers evolve. Therefore, we consider the animal population to provide fitness information as used by the evolutionary algorithms developing the robot programs.

Our work contrast with the use of robots to obtain or test models of animal behaviours Webb (2008). This line of research also uses EAs to obtain robot controllers Ponticorvo et al. (2007).

Why is evolving interactions with animal groups difficult?

Some of the primary difficulties that we face, as detailed further below, include: 1) Assessments of the collective response to a given stimulus pattern suffer from high variability. They are noisy both on a short timescale, i.e., within the multiple repeats used, and on longer timescales, i.e., repeats with different groups of animals can have drastically different fitness scores; 2) Robot heterogeneity also contributes to variability in fitness scores for identical stimulus patterns; and 3) A single fitness evaluation is expensive in time, meaning that we are unable to remedy these issues simply by obtaining larger samples.

Evolutionary computation has been applied successfully to problems with uncertainty in the fitness evaluation or in the longevity of a high-quality solution (Jin and Branke, 2005). One class of approaches to addressing problems with expensive fitness functions is approximation, including the approximation of fitness, problem, or evolutionary process (Jin, 2005). Here, either the cost of an evaluation is reduced – perhaps through a proxy, or the number of real fitness evaluations is reduced – for example, by taking an average of nearby points.

²Mixed societies are also designated bio-hybrid (Halloy et al., 2013) for integrating animals and robots.

Note, however, that each of these approaches depends on assumptions that make them less suitable for our approach. *Evolutionary approximation* assumes a fitness-distance correlation. We have not yet been able to establish a strong correlation between fitness and mutational space. *Problem approximation*. This is typically model-based (see also the “reality gap” in evolutionary robotics, Koos et al., 2013), and the specific science of discovery regarding improved/optimal communication patterns for guiding collective-level behaviours, via model-based surrogates, would require more fundamental understanding of the underlying system – and yet, that is precisely the knowledge that we seek.

One further prominent source of variability in fitness evaluation is that the hardware differs from unit to unit. This was observed in the application of evolutionary computation to programmable logic hardware (Thompson, 1996), in which the physical properties of a specific piece of hardware – rather than the designed-in properties – were exploited by the evolutionary algorithm. Specifically, the EAs would exploit a broader range of dynamics than the FPGA designers had intended. They are designed as digital computers, and the analog properties are intrinsic, but intended to be ‘engineered-out’ of the system. Consequently, solutions did not generalise well in other pieces of hardware. It was possible to find solutions with a generalised performance by evaluating in various FPGAs or by changing the region of FPGA used (Yao and Higuchi, 1999) (though in some cases at a fitness penalty).

Honeybees: relevant stimuli and behaviours

Communication is essential in animals, from social insects to humans. Honeybees have two types of complex communication signals which influence worker performance. The first type includes the waggle dance, tremble dance and grooming dance signals which influence certain workers. Such signals help to organize one or a few specific tasks within particular groups of workers, most of which respond to the same stimuli. The second type, “modulatory signals”, coordinate activity among different worker groups (Markl, 1985; Schneider and Lewis, 2004). Vibration signals of the honeybee, where the worker vibrates its body for 1–2 s, are examples of modulatory signals.

Collective behaviour describes the behaviour of a meta-organism, like a swarm, similar to the behaviour of an individual Bonabeau et al. (1999). Honeybee foraging behaviour is a very good example of natural “swarm intelligence”, because the group of foraging honeybees is not controlled by a central decision-making unit and because the collective decisions of colonies in varying environments have been found to act in an intelligent way (Schmickl et al., 2012).

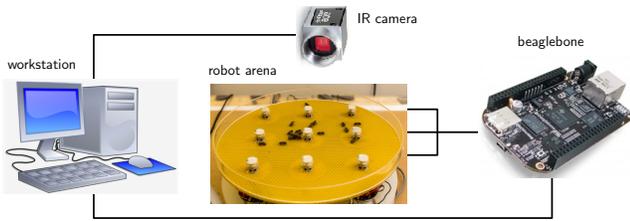


Figure 1: Schematic of the system that we are using to evolve a stopping vibration pattern.

Evolutionary problem

We are evolving robot controllers that interact with honeybees, here, with the specific aim of inducing aggregation behaviours. We have developed robots that produce a set of stimulus that honeybees respond to, namely vibration and temperature (Simpson and Cherry, 1969). The robots are also equipped with an air pump which is used to spread honeybees. The robots can sense their local environment through infra-red and vibration sensors. The arena consists of several robots in a lattice arrangement, and is observed by an overhead infra-red camera. The robots are controlled by a workstation where the evolutionary algorithm is run. This workstation also receives data from the infra-red camera. Figure 1 shows our experimental setup. We use young honeybees that are at most two days old, that do not fly. Therefore, the robot arena does not have a top cover.

The problem that we are trying to solve via evolutionary computation is to find a stopping vibration pattern. From previous work we know that honeybees respond to the queen piping signal by stopping (Simpson and Cherry, 1969). This signal is transmitted throughout the wax in the beehive (Michelsen et al., 1986b).

The vibration pattern that we are evolving comprises a single, repeated pulse. The pattern has a vibration period t_v followed by a pause period t_p . These periods are repeated as long as needed. The vibration period is characterised by a frequency f and an amplitude A . Depending on the experiment that we are conducting, either a subset or all of these genes are under evolutionary control.

In each experiment, one analysis of a vibration pattern comprises a sequence of actions. The available actions are vibrate, use the air pump, or do nothing. When there is more than one robot in the arena, only one is active while the others are passive. The implication is that the passive robot does nothing when the action in the sequence is to vibrate.

In order to compute the value of one action sequence we record the behaviour of the honeybees during the entire sequence of actions. The videos are processed as images and the basic functionality considers differences in pixels between pairs of frames from the video. Two pixels are considered different if their difference is greater than 25% of the maximum possible intensity.

Two measures are extracted using the following functions:

Background. The number of pixels that differ, in intensity above a pre-defined threshold, from an image of the arena without honeybees. This is a proxy for the number of honeybees that are in a particular region-of-interest. Inaccuracies in this estimate can occur when honeybees climb on top of each other or when they try to climb the walls, both of which would result in a lower background value.

Previous. The number of pixels that differ from a video frame 1 s before. This is a proxy for how fast honeybees are moving: a higher value indicates more movement in the population, either faster-moving or more honeybees moving.

The insights that we report on this paper were based on the following four experiments:

Experiment 1 12 honeybees were placed in a circular arena, and one robot was located in the centre of the arena. See Figure 2(a), which shows the setup and highlights the region of interest taken from a frame without honeybees. In contrast to later experiments, the chromosome here was limited to one gene: vibration frequency. The action sequence consisted of 30 s of vibration followed by 30 s of airflow in order to spread the honeybees in the arena. The value $e(\cdot)$ of one action sequence is defined as

$$e(c) = \sum_i \begin{cases} 1 & \text{if } \text{prv } i_a < T_p, \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

where c is the chromosome where the frequency is taken, i_a is the region of interest of the active robot in image i , prv is the **previous** function, and T_p is a threshold. The sum is only taken on the video images that correspond to the vibrate action. The range of this function is $\{0, 1, 2, \dots, 59\}$. The fitness value of a chromosome was the result of analysing three repetitions of the action sequence.³ A group of honeybees was used in at most 30 evaluation trials. The parameters of the EA were: population size 5; every individual in the population was subject to the mutation operator that consisted of adding Gaussian noise to the gene; the next population was selected among the best parents and offspring ($\mu + \lambda$ selection).

Experiment 2 12 honeybees were placed in an arena with stadium shape. Two robots were inside the arena, named active and passive. The arena image was divided into two regions of interest, also called active and passive, depending on the enclosing robot. See Figure 2(b) for an example of the regions of interest used by the image processing function. The chromosome was expanded to have

³We use the expression *evaluation trial/repetition* to refer to a repetition of an action sequence. This contrasts with *fitness evaluation*, which comprises several evaluation trials.

three genes: vibration frequency, pause period, vibration amplitude. The action sequence was similar to the previous experiment, but only the active robot vibrated. Both robots used airflow to spread the honeybees. The value of an action sequence is defined as

$$e(c) = \sum_i \begin{cases} 1 & \text{if } \text{prv } i_a < T_p, \\ 0 & \text{otherwise,} \end{cases} - \sum_i \begin{cases} 1 & \text{if } \text{prv } i_p < T_p, \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

where i_p is the region of interest of the passive robot in image i . Again the sum is only over the video images that correspond to the vibrate action. The range of this function is $\{-59, \dots, 0, 1, \dots, 59\}$. Again, the fitness value was the average of three repetitions of the action sequence. Each honeybee set was used in at most 30 evaluation trials. The EA used was as in experiment 1, except that the chromosome contained 3 genes. In each mutation event, one of these genes is selected randomly and modified by Gaussian noise.

Experiment 3 This experiment is similar to experiment 2, differing only by function $e(c)$ used to compute the fitness value of chromosome c , which is:

$$e(c) = \sum_i \begin{cases} \text{bkg } i_a & \text{if } \text{bkg } i_a > T_b \wedge \text{prv } i_a < T_p, \\ 0 & \text{otherwise,} \end{cases} - \sum_i \begin{cases} \text{bkg } i_p & \text{if } \text{bkg } i_p > T_b \wedge \text{prv } i_p < T_p, \\ 0 & \text{otherwise,} \end{cases} \quad (3)$$

where bkg is the function **background** described above, and T_b is a threshold. With this function we also take into account the presence of honeybees in a region of interest.

Experiment 4 This experiment does not involve an evolutionary algorithm but is instead an analysis of specific vibration patterns found by the EA in the previous three experiments. The main purpose is to check the quality of a vibration pattern. This experiment is characterised by the action sequence that the robots perform. In contrast with the previous experiments it may have more than one vibration action, each one with a different vibration pattern. Moreover, the length of each action can also differ.

Results and lessons learned

How should we interact with animal groups to maximise the information that they provide about the stimulus patterns presented, and specifically in the context of evolutionary search, *when can animal groups act as good fitness evaluators?* In this section we illustrate challenges and lessons



(a) Region of interest used in the experiment with the circular arena (b) Regions of interest used in the experiments with the stadium shaped arena.

Figure 2: Region of interest used in the experiments. These are images with a resolution of 600×600 pixels which is also the resolution of the videos collected during an action sequence. These images are a top view of the arena shown in figure 1.

learned with results from our experiments applying evolutionary computation to the problem of bio-hybrid interactions. We focus especially on issues relating to the animals providing the fitness function.

Animal habituation

Problem description When animals are presented a given stimulus for a long duration, their response may become diminished – that is to say, they become habituated to the stimulus and ignore it. Habituation is widely observed, e.g., the defensive reflex in *Aplysia*, (Castellucci et al., 1970). But we must pay attention to the specifics of the animal or society under investigation to ensure that the animals continue to provide good quality fitness evaluations.

Selecting an appropriate stimulus exposure period depends on several conflicting factors. A short period lowers risk of habituation, and additionally reduces overall time needed per experimental run. However, this must be balanced against a sufficient period to observe a response in behaviour – at the *group* level.

Experimental evidence In our experiments we presented vibration patterns for 30 s. We found that it was not uncommon for an otherwise successful vibration pattern to initially induce stopping of the honeybees, but before the end of the stimulus period the effect would wear off (see Figure 3, which illustrates four specific examples). We also analysed honeybee speed at different moments when the robot is vibrating, for all data from one run of an experiment. This showed that habituation occurs systematically: there is more similarity between what happens during the first 10 frames and the second 20 frames (p-value 3.9×10^{-1} Mann-Whitney U-test), than between the first 10 frames and the last 10 frames (p-value 5.1×10^{-23}).

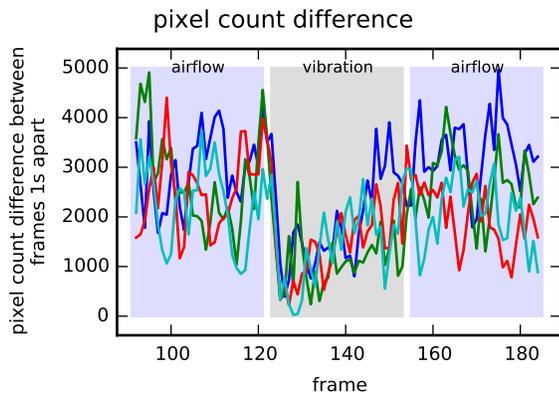


Figure 3: Plot of the **previous** function before and after vibration from four videos recorded during experiment 4. Before vibration honeybees move at a regular speed. When the vibration starts, they react by almost stopping, but then they resume their normal speed when the robot is still vibrating.

Lesson learned Interestingly, since evolutionary optimisation requires fitness gradients in order to select between different candidate chromosomes, observing habituation does not *inherently* indicate an over-length stimulus period; but if all chromosomes exhibit significant habituation then reducing the stimulus period may offer a more efficient balance. There is also opportunity for this period to co-evolve, with the aim of providing fitness gradients for longer.

Animal fatigue

Problem description How long can one animal group continue to provide accurate and useful fitness evaluation? We have two general motivations to use a single group of animals more than once. Firstly, the animals are a limited resource – in our case, the number of juvenile honeybees available depends on factors such as the phase of the season and recent weather. Secondly, there is a time penalty associated with each change of in the experimental arena. In our case, it takes approximately 4 min to count honeybees, to give them a “familiarisation” period where their stress levels from moving environment are allowed to reduce, and to remove honeybees from the arena; this is in comparison to 1 min per evaluation trial/repetition. On the other hand, animals need to rest and to get food, otherwise they will get tired and their activity levels drop substantially. In our experiments, bees are not fed when they are in the arena.

Overall, we have a tradeoff that relates to maximal “yield” of evaluations per animal group. Note that this issue differs from habituation since it refers to a change in behaviour *between* multiple trials, rather than *within* one trial.

Experimental evidence Prior work indicates that juvenile honeybees can sustain at least 30 minutes in experiments of

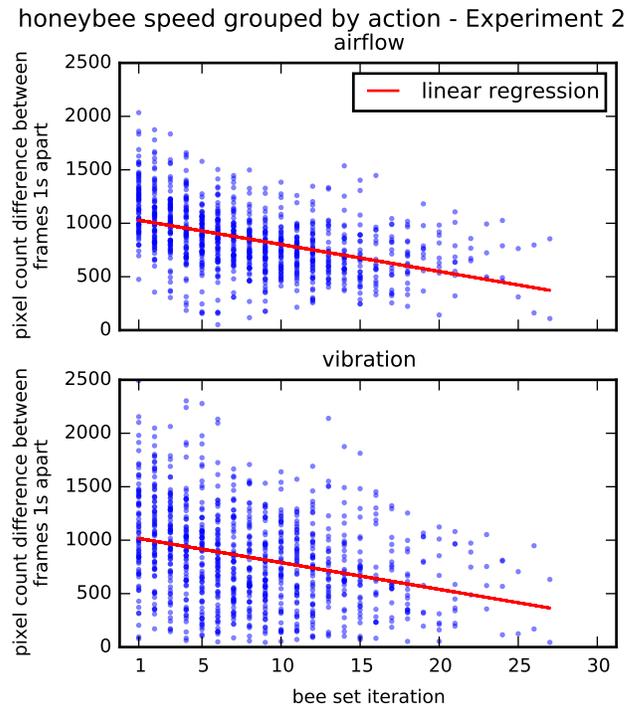


Figure 4: Honeybee speed during an evaluation trial grouped by the action versus the “time” the trial was done with a honeybee set. The data is from experiment 2. Recall that the action sequence was 30 s of vibration followed by 30 s of airflow.

this form (see e.g., Szopek et al., 2013). If we exchange the animals between each evaluation trial, only 20% of the experimental time would provide evaluations; conversely, retaining one group for 30 evaluation trials would mean 88% of the time is used for evaluations.

Figure 4 shows the honeybee speed (the value of function *prv* applied to an entire image) versus honeybee set iteration, grouped by the action in the evaluation trial. The plots also contain the linear regression. As can be seen, the longer honeybees are in the arena, the slower they move, which is a sign of fatigue. This contrasts with animal habituation, where animals stop reacting to the signal. In our case, they resume their normal, faster behaviour (contrast with Figure 3).

One of the consequences of animal fatigue is that we changed honeybees earlier than 30 min in the majority of experiments. The mode of the number of evaluation trials per honeybee set was 12 in experiments 2 and 3. If the human operator that was overseeing the experiment noticed any considerable number of stopped honeybees, they would interrupt the process to refresh the honeybees.

Lesson learned This result suggests that the specific environmental conditions may have a large impact on how long

a group of honeybees is able to sustain activity. Since we aim for an overall evolutionary process with a high level of autonomy, any subjective decision such as when the animals are too fatigued to provide accurate fitness evaluations should be avoided where possible. By limiting the time that one group is used to 12 minutes, we would substantially reduce the cases where a subjective decision was required. A more formal method to assess animal state could be used, depending on their availability, see e.g., Cazenille et al. (2015).

Heterogeneity of animal response

Problem description One of the challenges of working with animals is the diversity of their behavioural repertoire. Working with young animals can reduce this breadth somewhat, and some physical constraints can also keep animals to perform the behaviours we wish to study. Nonetheless, our pursuit to use animal groups to evaluate stimulus fitness depends on a certain degree of repeatability in behavioural response, and animals are not deterministic mathematical functions. For instance, our lab has found evidence of multiple types in the locomotion behaviours of the juvenile honeybee (Kengyel et al., 2015).

Depending on the animal under study, it may be possible to pre-select individuals to reduce the heterogeneity, e.g., on gender if they present sexual dimorphism; or to pre-select group composition, e.g., if social hierarchy is well understood. However, in the general case this may be a scientific endeavour in its own right. It is of course very common in natural sciences to reduce uncertainty through multiple replicates. The increased cost in time from the number of replicates/trials used for each fitness evaluation must be justified carefully.

Experimental evidence As an example of animal behaviour heterogeneity, Figure 5 shows overviews of honeybee behaviour from runs of experiment 4, in which we analysed chromosomes that were found in experiment 1. The action sequence consisted of 2 min of airflow, followed by 30 s of vibration, followed by 30 s of airflow. The frequency of the vibration pattern depends on the chromosome. The other parameters of the vibration pattern are: vibration period 1 s, pause period 0.1 s, amplitude 100%.⁴ The overview uses the two functions **background** and **previous** to process the recorded videos. All videos were divided in segments of 30 s, as represented in the vertical axis. All honeybee sets were used only once (B1–B10 on the horizontal axis). For each video segment of a honeybee set, we computed the average of the **background** value, which is represented by the bar height, and the average of the **previous** value, which is represented by the bar colour. For clarity we only focused on the 30s previous to vibration and in the video segment with

⁴Amplitude unit is relative to the maximum permitted by the robots.

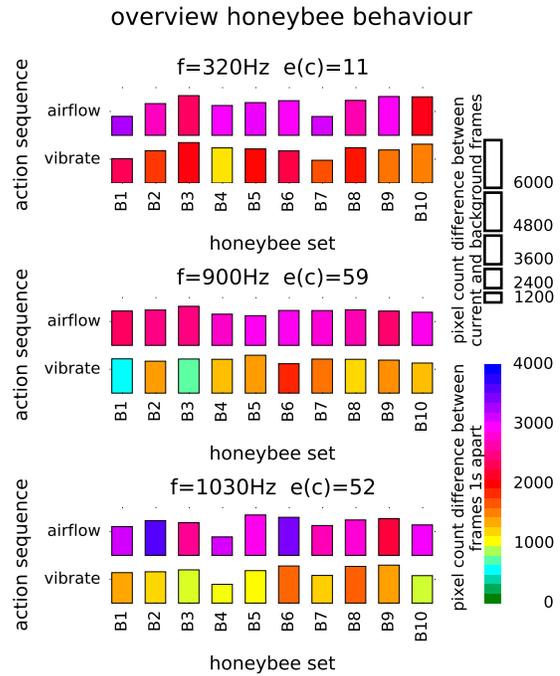


Figure 5: An overview of honeybee behaviour from several runs of experiment 4 using the vibration frequency of chromosomes from experiment 1. For each video segment of a honeybee set we computed the average of the *background* value which is represented by the bar height and average of the *previous* value which is represented by the bar colour. See text for explanation.

vibration. As can be seen, the honeybees move slower in the segments where there is vibration than in the segments with airflow. Moreover, the difference is higher when the fitness value is higher.

Lesson learned In experiment 4 we obtained a high number of replicates for several specific chromosomes, to quantify the number of replicates required to obtain an accurate estimate of the underlying quality (data not shown). Although we are using juvenile honeybees that have a small behaviour repertoire (Kengyel et al., 2015), we nevertheless obtained different behaviours for the same stimuli. At this point we have not been able to fully isolate the source of this variability. We have run several clustering algorithms on the data shown on these figures, but the algorithm did not produce a sharp distinction between segments of different stimuli.

In order to solve the problem of behaviour heterogeneity, we can increase the number of repetitions done per fitness computation. This has an added cost of experimental time. Besides increasing the number of repetitions, there are var-

ious potential strategies such as excluding outlier replicates from the fitness calculation. These could be based on the observed distribution, e.g. maximum and minimum values, or on detected characteristics that correspond to specific issues such as fatigue. Although such exclusions necessitate throwing out experimental time, it is an option to be considered, especially if the animals show signs of fatigue or habituation during one action sequence.

The results in Figure 5 show that there is reasonable consistency in the honeybee behaviour among the multiple replicates under the same stimuli patterns. Moreover, the chromosomes deemed to be high (low) fitness from three evaluation trials are confirmed to exhibit generally low (high) movement, indicating that under some circumstances three trials are enough for an evolutionary method to succeed.

Chromosome Evaluation

Above, we discussed problems that may occur when encountering variability in animal behaviour. Here we consider issues relating to the robot control and characteristics, and how they may influence chromosome evaluation.

When we want to evolve a robot controller we have to take into consideration the possibility of the EA exploiting unforeseen and idiosyncratic properties of the robot, as discussed (see Thompson, 1996). One possibility to mitigate this problem is to evaluate the chromosome in different robots, aiming to avoid the local overfitting. In our case this corresponds to employing multiple arenas, either sequentially or simultaneously. In the latter case, each arena would have a set of honeybees, and to evaluate a new chromosome, one arena would be chosen randomly and within that arena, the role of each robot is also chosen randomly. This scheme would also have to be flexible enough to handle the animal fatigue, i.e., if animals in one arena were detected to be tired, that arena should be avoided until it is replenished with fresh animals.

When a new population of offspring has to be evaluated, we opt to evaluate each chromosome in a row. That is to say, the first N action sequences are for evaluating chromosome c_1 , then the second N action sequences are for chromosome c_2 , up to the last chromosome. In order to combat possible habituation to the same stimulus pattern, several actions can be performed. Each subsequent action sequence could have segments without any stimuli or with an alternative stimulus to spread the animals. We have opted to have a segment with vibration stimulus followed by segment with airflow, to inject “noise” into the animal population. As an alternative scheme, we could instead intermix the evaluation of chromosome c_1 with the evaluation of other chromosomes, to partially mitigate overfitting of a specific honeybee set to one chromosome.

All of these proposals to chromosome evaluation add extra layers of complexity to the software, with the required

time to debug it; though once validated, the experimental runtime would not be substantially impacted.

Conclusion

The application of evolutionary computation to real-world tasks often introduces difficulties not encountered by theoreticians. However, this first foray into using animals other than humans to provide fitness evaluations, has introduced an array of interesting challenges as well as annoying ones. This paper has discussed several specific aspects that require special attention when using animal groups as part of the fitness function. Many of the issues could be straightforwardly solved by enlarged sample sizes if time and animal supplies were unlimited, but this being counterfactual necessitates tradeoffs between accuracy and experimental throughput.

Combined with animal resting time, a single run of the evolutionary algorithm with a reasonable number of generations may take a long time. In experiments 1 to 3, population size was five and the number of evaluations done per chromosome was three. Since each chromosome evaluation took about 1 minute, ten generations took less than one working day. While ten generations are typically insufficient for evolutionary methods to find optimal solutions to complex problems, it is nonetheless capable of yielding substantial improvement in chromosome quality.

The most obvious resolution to many of the issues highlighted above come with a tradeoff against time – an extremely precious resource – which makes it crucial to minimise the extensions used and to seek more sophisticated resolutions. Parallelisation is a powerful method to increase experimental throughput, and although we must be careful to ensure isolation from one repeat to the next, offers much promise in the application of evolutionary computation to bio-hybrid system design.

Acknowledgements

This work is supported by EU-ICT project ASSISIBf no. 601074.

References

- Bonabeau, E., Dorigo, M., and Theraulaz, G. (1999). *Swarm Intelligence: From Natural to Artificial Systems*. Oxford.
- Bonnet, F., Kato, Y., Halloy, J., and Mondada, F. (2016). Infiltrating the zebrafish swarm: design, implementation and experimental tests of a miniature robotic fish lure for fish–robot interaction studies. *Artificial Life and Robotics*, 21(3):239–246.
- Caprari, G., Colot, A., Siegart, R., Halloy, J., and Deneubourg, J.-L. (2005). Animal and robot mixed societies: building cooperation between microrobots and cockroaches. *IEEE Robotics & Automation Magazine*, 12(2):58–65.
- Castellucci, V., Pinsker, H., Kupfermann, I., and Kandel, E. R. (1970). Neuronal mechanisms of habituation and dishabituation of the gill-withdrawal reflex in aplysia. *Science*, 167(3926):1745–1748.

- Cazenille, L., Bredeche, N., and Halloy, J. (2015). *Multi-objective Optimization of Multi-level Models for Controlling Animal Collective Behavior with Robots*, pages 379–390. Springer.
- Cully, A., Clune, J., Tarapore, D., and Mouret, J.-B. (2015). Robots that can adapt like animals. *Nature*, 521(7553):503–507.
- Dawkins, R. (1986). *The Blind Watchmaker*. Longman Scientific and Technical.
- Donati, E., Worm, M., Mintchev, S., van der Wiel, M., Benelli, G., von der Emde, G., and Stefanini, C. (2016). Investigation of collective behaviour and electrocommunication in the weakly electric fish, mormyrus rume, through a biomimetic robotic dummy fish. *Bioinspiration & Biomimetics*, 11(6):066009.
- Funes, P., Sklar, E., Juillé, H., and Pollack, J. (1998). Animal-animat coevolution: Using the animal population as fitness function. *From Animals to Animats*, 5:525–533.
- Goumopoulos, C., Christopoulou, E., Drossos, N., and Kameas, A. (2004). *The PLANTS System: Enabling Mixed Societies of Communicating Plants and Artefacts*, pages 184–195. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Halloy, J., Mondada, F., Kernbach, S., and Schmickl, T. (2013). Towards bio-hybrid systems made of social animals and robots. In *Conference on Biomimetic and Biohybrid Systems*, pages 384–386. Springer.
- Jin, Y. (2005). A comprehensive survey of fitness approximation in evolutionary computation. *Soft computing*, 9(1):3–12.
- Jin, Y. and Branke, J. (2005). Evolutionary optimization in uncertain environments—a survey. *IEEE Transactions on evolutionary computation*, 9(3):303–317.
- Kengyel, D., Hamann, H., Zahadat, P., Radspieler, G., Wotawa, F., and Schmickl, T. (2015). Potential of heterogeneity in collective behaviors: A case study on heterogeneous swarms. In *International Conference on Principles and Practice of Multi-Agent Systems*, pages 201–217. Springer.
- Koos, S., Mouret, J.-B., and Doncieux, S. (2013). The transferability approach: Crossing the reality gap in evolutionary robotics. *IEEE Transactions on Evolutionary Computation*, 17(1):122–145.
- Landgraf, T., Bierbach, D., Nguyen, H., Muggelberg, N., Romanczuk, P., and Krause, J. (2016). Robofish: increased acceptance of interactive robotic fish with realistic eyes and natural motion patterns by live trinidadian guppies. *Bioinspiration & biomimetics*, 11(1).
- Landgraf, T., Rojas, R., Nguyen, H., Kriegel, F., and Stettin, K. (2011). Analysis of the waggle dance motion of honeybees for the design of a biomimetic honeybee robot. *PloS one*, 6(8):e21354.
- Lorenz, K. (1950). Ganzheit und teil in der tierischen und menschlichen gemeinschaft. In *Studium generale*, volume 3, pages 455–499. Springer.
- Markl, H. (1985). Manipulation, modulation, information, cognition: some of the riddles of communication.
- Michelsen, A., Kirchner, W. H., Andersen, B. B., and Lindauer, M. (1986a). The tooting and quacking vibration signals of honeybee queens: a quantitative analysis. *Journal of Comparative Physiology A*, 158(5):605–611.
- Michelsen, A., Kirchner, W. H., and Lindauer, M. (1986b). Sound and vibrational signals in the dance language of the honeybee, *apis mellifera*. *Behavioral ecology and sociobiology*, 18(3):207–212.
- Ponticorvo, M., Walker, R., and Miglino, O. (2007). Evolutionary robotics as a tool to investigate spatial cognition in artificial and natural systems. In Loula, A., Gudwin, R., and Queiroz, J., editors, *Artificial Cognition Systems*, pages 210–237.
- Schmickl, T., Bogdan, S., Correia, L., Kernbach, S., Mondada, F., Bodi, M., Gribovskiy, A., Hahshold, S., Miklic, D., Szopek, M., et al. (2013a). Assisi: Mixing animals with robots in a hybrid society. In *Biomimetic and Biohybrid Systems*, pages 441–443. Springer.
- Schmickl, T., Szopek, M., Bodi, M., Hahshold, S., Radspieler, G., Thenius, R., Bogdan, S., Miklic, D., Griparic, K., Haus, T., et al. (2013b). Assisi: Charged hot bees shakin’ in the spotlight. In *Procs. Self-Adaptive and Self-Organizing Systems (SASO)*, pages 259–260. IEEE.
- Schmickl, T., Thenius, R., and Crailsheim, K. (2012). Swarm-intelligent foraging in honeybees: benefits and costs of task-partitioning and environmental fluctuations. *Neural Computing and Applications*, 21(2):251–268.
- Schneider, S. and Lewis, L. (2004). The vibration signal, modulatory communication and the organization of labor in honey bees, *apis mellifera*. *Apidologie*, 35(2):117–131.
- Seeley, T. D. (1995). *The wisdom of the hive: the social physiology of honey bee colonies*. Harvard University Press.
- Simpson, J. and Cherry, S. M. (1969). Queen confinement, queen piping and swarming in *apis mellifera* colonies. *Animal Behaviour*, 17, Part 2:271 – 278.
- Szopek, M., Schmickl, T., Thenius, R., Radspieler, G., and Crailsheim, K. (2013). Dynamics of collective decision making of honeybees in complex temperature fields. *PLOS ONE*, 8(10):1–11.
- Takagi, H. (2001). Interactive evolutionary computation: Fusion of the capabilities of ec optimization and human evaluation. *Proceedings of the IEEE*, 89(9):1275–1296.
- Thompson, A. (1996). An evolved circuit, intrinsic in silicon, entwined with physics. In *International Conference on Evolvable Systems*, pages 390–405. Springer.
- Tinbergen, N. (1948). Social releasers and the experimental method required for their study. *The Wilson Bulletin*, 60(1):6–51.
- Webb, B. (2008). Using robots to understand animal behavior. In *Advances in the Study of Behavior*, volume 38, pages 1–58. Academic Press.
- Yao, X. and Higuchi, T. (1999). Promises and challenges of evolvable hardware. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 29(1):87–97.