

Evolving robot controllers for a bio-hybrid system

Pedro Mariano¹, Ziad Salem², Rob Mills¹,
Stefan Schönwetter-Fuchs-Schistek², 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, 8010 Graz, Austria
plmariano@ciencias.ulisboa.pt

Abstract

In this paper we report the first results of evolving bio-hybrid societies. Our goal is to have robots that are integrated in an animal society, and here we evolve robot controllers using animals as fitness providers, directly judging the success of integration. In particular, we are using juvenile honeybees and robots that are able to produce vibration patterns. Previous studies have shown that honeybees react to different vibration patterns, such as exhibiting freezing or stopping behaviours. In this paper we investigate whether we are able to evolve a vibration pattern that acts as a locally acting ‘stop signal’ for bees. Honeybees were placed in two containers with no communication between them: one with an active, vibrating robot, and a second with a passive robot. Post-hoc evaluations of key evolved digital genotypes generally confirm fitness values obtained during evolution. We also tested the transferability of key genotypes to a single container, in which bees are free to visit one vibrating and two dummy robots. Encouragingly, most genotypes are able to selectively stop bees, i.e., only in the vicinity of the vibrating robot, despite having been evolved under the more constrained setup. These results speak to the value of an evolutionary approach for discovering how to interact with animals.

Introduction

Information exchange among animals is crucial to social behaviour. The proximate mechanisms that organisms have developed to exchange information show a high diversity. Besides their varied physical implementations, they differ concerning specific aspects: be it via direct or indirect communication, the latter mainly in the form of ‘stigmergic’ interaction, prominently studied in social insects (Deneubourg et al., 1990). There are many more direct forms of communication in animals. Within honeybees, which are focal in our research, multiple direct communication signals have been identified: various dances (Anderson and Ratnieks, 1999) include the waggle dance, which recruits other bees to food sources (Seeley, 1994); as well as trophallactic food exchange among bees (Camazine et al., 1998; Schmickl and Karsai, 2016). Bees exhibit a characteristic stopping behaviour to the queen piping signal (Simpson and Cherry, 1969), which is transmitted throughout the wax in the beehive (Michelsen et al., 1986). Recently, short-pulsed vibra-

tions have been associated with bee-to-bee encounters (collision) in the beehive (Ramsey et al., 2017). These short-term vibration signals also seem to have an effect on the motion pattern exhibited by bees (Mariano et al., 2017) but is not yet fully understood, as also the role of the honeybee comb itself as a vibrational information centre is not yet fully deciphered (Bencsik et al., 2015).

Our EU project “ASSISIBf” (Schmickl et al., 2013) is in the area of bio-hybrid societies, in which we use robotic devices to interact with animal societies in a closed-loop. We aim for a method to discover some aspects of “the language of the animals”, i.e., those stimuli patterns that animals will respond to, using two animal species, zebrafish and honeybees. It is our aim to facilitate such interaction automatically, identifying robotic behaviour parameters through experimentation that is guided by evolutionary search. In this context, gathering data is very expensive in comparison to typical applications in evolutionary computation (EC). Working with animals requires an adequate time frame, and no technical advance will change the speed of the animal behaviour. Although there of course exist other EC applications with expensive fitness evaluation, typical approaches of modelling via surrogates is not feasible when we are exploring a behaviour that is itself not well understood, and thus we are not able to employ modelling to leverage the search.

While we knew bees stop or at least slow down significantly with specific vibration patterns produced by other bees in the hive we wanted to investigate if we can use evolutionary computation to let robots learn effective variants of such vibration patterns in an autonomous way. In order to build an “animal language” we evolved robot controllers to produce vibration pulses that cause juvenile honeybees to stop. We expect that this behaviour can be further used as a building block of more complex behaviours.

This work is situated in the area of evolutionary robotics (ER) (Harvey et al., 1997; Nolfi and Floreano, 2000; Silva et al., 2016), which is an evolution-inspired technique, used to generate robot behaviours that are difficult to derive analytically from the robot’s mechanics and task environment

(Bongard, 2008). Conventional ER has diverse applications including controllers of robots that resemble or mimic animals (e.g., Cully et al., 2015). We use online ER to discover how to control robots that interact with groups of real animals, in an original approach where the animals are the fitness providers, guiding the evolutionary process. In online ER the robots' control evolves while robots perform their behaviour in the real context in which they operate, counter to offline ER where robot controllers are evolved in simulation.

Our recent study (Mariano et al., 2017) identified the main new issues raised by ER with evolution guided by the animals. However, that work did not include specific results of the evolutionary process. This paper presents, for the first time, results of online ER interacting with animals with the latter providing the fitness that guides the evolutionary process. In addition, we test the evolved digital genotypes that define the robot programs in a less-constrained environment, in order to study their generalisation capability.

In the next section we present related work and then we have a section detailing the problem of evolving vibration patterns to stop juvenile bees. In the following section the evolution results are presented as well as a detailed analysis of their variability, based on re-assessment of selected patterns. The next section shows how well evolved vibration patterns transfer to a different but related task. The paper closes with a section of discussion and conclusions.

Related Work

As technology becomes increasingly pervasive in the human world, so it does in the animal kingdom. Besides improved instrumentation and telemetry devices (e.g. Begg et al., 2005), scientific studies use artifacts to stimulate animals (Reaney et al., 2008; Bonnet et al., 2018), with an increasing sophistication. Beyond using robots to emit stimuli repeatedly, bio-hybrid systems comprise animals and robots that each influence one another, i.e., the interactions are closed-loop and flow both from robot to animal and from animal to robot. A major milestone is the robotic cockroaches of Halloy et al. (2007), in which a group of robots were able to successfully integrate with a group of cockroaches, and, when programmed with unusual environmental preferences were able to steer the natural cockroaches into a different decision than they would take alone. Closed-loop coupling of animals and robots have been explored with various organisms such as ducks (Vaughan et al., 2000), fish (Swain et al., 2012; Bonnet et al., 2018), and honeybees (Landgraf et al., 2012; Griparic et al., 2017; Stefanec et al., 2017).

Unlike all of these works that use an *a priori* model of the animal behaviour to guide the robot design, our present work attempts to develop the means of robot-animal interaction at runtime: Specifically, we use the animals as fitness providers in an evolutionary algorithm, an innovation not seen before in bio-hybrid systems. Arguably the closest works, conceptually, involve humans interacting with agents of some kind,




Figure 1: Schematic of the closed-loop bio-hybrid system that we use to evolve a stopping vibration pattern. The vibrations stimulate the animals, and observations of their behaviour influences characteristics of the future vibrations.

(e.g., Funes et al., 1998; Dawkins, 1986). However, to our knowledge animals other than humans have not previously been used as part of an evolutionary algorithm.

In terms of fitness measures obtained from a collective in guided EC we can trace back to Zhang and Cho (1999), where a fitness value is obtained by the sum of individual properties of the agents forming the group. A study worthy of note here is that of Li et al. (2013), which uses an adversary co-evolutionary algorithm to automatically parameterise behavioural models of (simulated) animals. A population of models evolves under a pressure of detection from a population of classifiers, whose aim is to tell apart model-derived data samples and true animal behaviour samples. This study uses ER in the pursuit of understanding animal behaviour through interaction of real animals with robots. In our work we observe animals without identifying them, therefore we have an aggregating measure of the group behaviour, based on approximate counts of the number of bees around each robot.

Evolution of Localised Vibration Pattern

Experiment 1 setting

We are evolving robot controllers that interact with honeybees, here, with the specific aim of inducing aggregation behaviours. We perform experiments in a dark environment so to remove the influence of visual cues on behaviour, instead focusing on the stimuli emitted by the robots. Moreover, we use juvenile bees, with a maximum age of 48h, which do not fly and accordingly can be stimulated through the arena floor. We have developed immobile robots that produce a set of stimuli that honeybees respond to, namely vibration and temperature (Griparic et al., 2017). The robots are also equipped with an airflow stimulus, to spread honeybees; and an LED for debugging. For feedback purposes we use an

infra-red camera (Basler acA2040-25gmNIR) mounted approx. 1.2m above the arena in order to view the whole area where robots and honeybees are. The robots are arranged in lattice with a separation of 9cm, and are controlled by “beaglebone” single-board computers. In addition, there is a workstation that receives data from the infra-red camera and runs the evolutionary algorithm. Figure 1 shows the experimental setup that we are using.

To measure honeybee behaviour (movement) we analyse the videos immediately after they are recorded. Before placing a new set of honeybees in the arena, we record an image of the arena that we will call *background*. The analysis is restricted to a set of regions of interest (ROIs), R , that depends on the experiment. For each video frame number t , F_t , and ROI r we computed the following four values:

number of bees, NB_r^t : number of pixels whose intensity is different between frame F_t and the background image B . Two pixels are considered similar if the difference, in percentage, of their intensities is lower than threshold SCT .

total bee movement, TBM_r^t : number of pixels whose intensity is different between frames F_t and $F_{t-\Delta_f}$. This is a proxy for bee speed as we are not tracking individual bees, and we are not computing displacement. If we used frames further apart, we would get similar values. The threshold SCT is also used to compute whether two pixels are different or similar.

bee acceleration, BA_r^t : difference between bee movement recorded in frames F_t and $F_{t-\Delta_v}$.

average bee movement, ABM_r^t : total bee movement in frame F_t divided by the difference between number of bees in frames F_t and $F_{t-\Delta_f}$. There are experiments where there is no physical separation in the arena zones that correspond to the ROIs, meaning that honeybees can move out or into a region of interest. ABM compensates for this characteristic (cf. sensitivity of TBM).

Preliminary evaluation of robots

Before attempting to evolve robot controllers, we performed two pilot studies to test our robots’ ability to influence bee behaviour. This section reports very briefly on the study, to better contextualise the ER work later in the paper.

To this end, we placed groups of honeybees in a circular arena with a single robot. The robot played a sequence of vibration patterns separated by intervals with no stimuli. Each action sequence $S = (a_1, a_2, \dots, a_7)$ was composed of actions¹ $a_{2n-1} = N$ and $a_{2n} = V$ for 30 s each. Vibration actions used maximum amplitude and were pulsed on/off with a frequency of f Hz for a vibration period $t_v = 0.9$ s and a pause period of $t_p = 0.1$ s; prior studies using pure-tones had

¹ In the rest of the paper we will denote a no stimuli, airflow or vibration action/segment by the letters N, A and V, respectively.

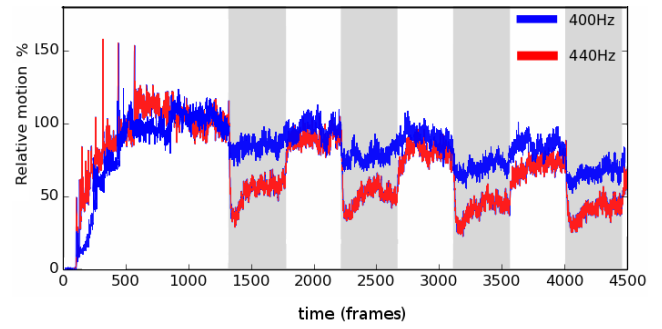


Figure 2: Mean bee motion during pulsed-vibration experiments (grey bands: vibration emitted; white bands: no stimulus), for groups of 15 bees, $f = 400$ Hz ($n=9$), $f = 440$ Hz ($n=8$). The bee motion is reduced during the vibrating periods, with an effect size that depends strongly on frequency.

insignificant effects. We performed experiments with groups of 1, 5, and 15 honeybees, and with f of 300, 400, 440, and 500 Hz, at least $n = 8$ replicates per condition.

We recorded the behaviour of honeybees as mentioned in the previous section. We then analysed the videos and we measured bee motion in a single circular ROI. Figure 2 shows an estimate of bee motion for two frequencies, relative to the pre-vibration motion rate, averaged over all replicates. As can be seen bee speed drops drastically at the onset of each vibration period, corresponding to bees stopping.

We can summarise the findings as follows: (1) The vibrating robots can indeed significantly affect bee behaviour; but pulsing vibrations appeared to be crucial: $t_p = 0$ (pure-tone) is ineffective but $t_p = 0.1$ can be very effective. (2) The behaviour shows sensitivity to frequency (even a 10% change in f , Fig. 2). (3) The effect was greater with more bees, indicating a socially-mediated response (not depicted).

We performed another pilot study to evaluate the range at which vibrations are felt, in which we placed bees in an arena whose centre is 9, 18, 27 cm away from the vibrating source ($f = 440$ Hz, $t_p = 0.1$). We observed that bee behaviour is modulated up to 18 cm away.

Armed with this knowledge, we set out to explore how the behavioural modulation could be improved, across a broader range of the parameters shown to have an influence; and moreover, to exploit the modulation for a novel task: whether the robots could induce localised stopping. Given these goals, we were interested in pursuing our exploration through an automated process.

Evolutionary problem

The problem that we are trying to solve via evolutionary computation is to optimise a *localised* stopping vibration pattern for juvenile honeybees. We use the above-mentioned pattern with a vibration period t_v followed by a pause period t_p . These periods are repeated as long as needed. The




Figure 3: Example of a video frame from an experiment with two circular arenas, and 12 bees in each arena.

vibration period is characterised by a frequency f and an amplitude v . In the evolutionary experiment we have set $t_v + t_p = 1$ s. We thus have three vibration parameters under evolutionary control: f , v , and t_p (we arbitrarily chose t_p as the independent variable in the previous equation). The domain of gene f is $\{300, 310, 320, \dots, 1500\}$ in Hz. This gene was mutated by adding Gaussian noise with mean zero and standard deviation 120,

$$f' = \left(\left\lceil \frac{f + \mathcal{N}(0, 120)}{10} \right\rceil 10 - 300 \right) \bmod 1200 + 300.$$

If the value exceeds either the minimum or maximum value, we wrap around. The domain of gene t_p is $\{100, 110, 120, \dots, 900\}$ in millisecond. The minimum value of 100 ms was chosen due to hardware limitations in generating patterns with a shorter period. This gene was mutated by applying the equation

$$t_p' = \left(\left\lceil \frac{t_p + \mathcal{N}(0, 300)}{10} \right\rceil 10 - 100 \right) \bmod 900 + 100.$$

The domain of gene v is $\{5, 10, 15, \dots, 50\}$, normalised so that maximum amplitude corresponds to 50. This gene was mutated by applying the equation

$$v' = \left(\left\lceil \frac{v + \mathcal{N}(0, 10)}{10} \right\rceil 10 - 5 \right) \bmod 45 + 5.$$

The setup used consisted of two circular arenas, each with a robot in the centre. One robot was active, meaning it played a vibration pattern, while the other robot, named passive, did nothing. We recorded a video of the bee behaviour, as depicted in Figure 3, together with the ROIs used.

We used a $(\mu + \lambda)$ evolutionary strategy (ES) with population size 5. In each generation, every genotype had one randomly chosen gene mutated. The offspring were evaluated and the next generation was composed of the best 5 individuals from the set of parents and offspring. We performed 10 generations in each ES run.

Each genotype was subjected to 3 evaluations and its fitness value was the average. Each evaluation consisted of an action sequence with 30 s of vibration followed by 30 s of airflow. This airflow action was used to disperse honeybees aggregated due to the previous vibration pattern. In each generation, the evaluation order of all offspring was

randomised, for instance we could perform first evaluation of genotype 1, then first evaluation of genotype 3, then second evaluation of genotype 1, and so on. For any genotype evaluation we only considered the frames F_t of the vibration segment. The computed value was:

$$\sum_t \begin{cases} 1 & \text{if } TBM_t^A < T \\ 0 & \text{otherwise} \end{cases} + \begin{cases} -1 & \text{if } TBM_t^P < T \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where the superscripts A and P represent the active and passive ROIs respectively, and T is a constant threshold. To compute TBM we set $SCT = 5\%$ and $\Delta_f = 2$. A total of 58 frames were used (the frame rate was 2Hz).² The domain of the fitness value is $[-58, 58]$.

Each arena had twelve honeybees. Each bee set was used in exactly 12 genotype evaluations (see Mariano et al., 2017). When we changed the set of honeybees, we also replaced the wax floor to avoid any pheromones from the previous honeybee set affecting the behaviour of the new ones. We also waited 30 s before evaluating any genotypes with the new honeybee set, to reduce honeybee stress levels from being moved from the bee keeping box to the arena. To have independent experiments, replaced bees are not used in further evaluations (they are returned to the hive).

Experiment 1 results

Figure 4 shows the fitness values over time in each run. In all runs, the mean fitness value of the initial population was around zero. From seven runs, four produced high quality genotypes (runs 3, 4, 6 and 7), in the sense of obtaining genotypes in the top twenty with fitness above the 80th percentile. Moreover, the best genotype appeared in generation 1 of run 6.

There were also runs where there was a small initial improvement, but then the fitness value stabilised (runs 1, 2 and 5). Figure 5 shows the histogram of genotype evaluations. There are over 300 genotypes that have an evaluation of zero. For this reason we truncated the vertical axis in order to focus on other counts of fitness value. The figure also shows how many genotypes are in the top 5% and bottom 40% of the fitness domain.

To give an impression of the fitness landscape, Fig. 6 shows genotypes that obtained a fitness value in the top 5% or bottom 40% of the fitness domain. As can be seen, there is not a clear separation between these extreme sets. From these plots we can observe that the genotypes in the top are concentrated around ($f = 878.8$ Hz, $v = 35.8\%$, $t_p = 340$ ms), although only amplitude is sharply defined. Further analysis of the evolutionary paths in genotype space shows that the runs with small improvement started away from this point and evolution failed to find this area.

²The total bee movement value depends on frames where there was no vibration and thus were excluded.

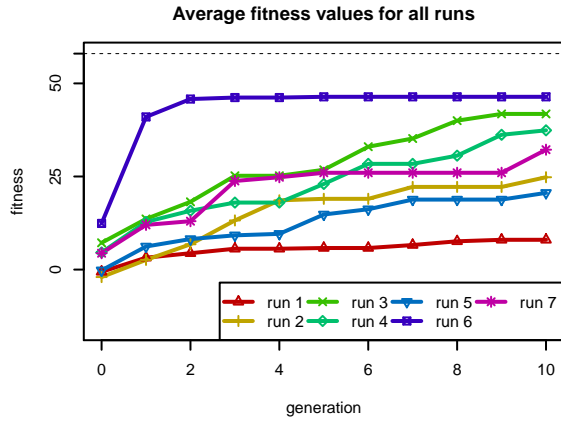


Figure 4: Mean fitness over time in the seven runs performed. Maximum fitness value is 58. Typically the ES is able to improve the efficacy of the vibration patterns, but in some runs the performance is meagre, suggesting that the fitness landscape does not always provide a clear gradient.

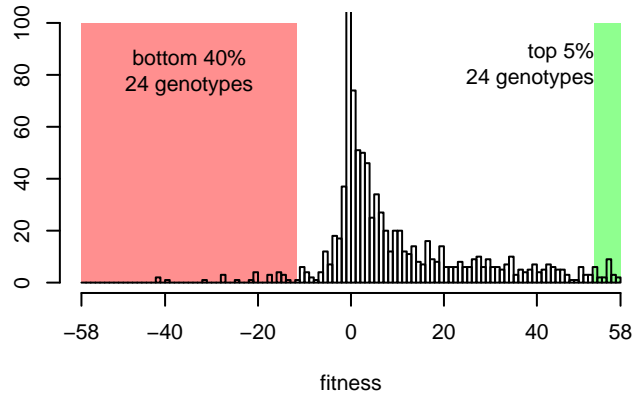


Figure 5: Histogram of genotype evaluations. We focused on fitness values different from zero. There are 341 genotypes that had a fitness value of zero.

Post-hoc Analysis of Evolved Genotypes

In order to quantify the improvement in fitness provided by the evolutionary process, we analysed a sample of nine genotypes in further detail. These were the top three fitness values overall, C_T , and two genotypes from each of the three runs with highest ranking genotypes, in the following way: lowest fitness value in the last generation, C_L ; and highest fitness value in the first generation, C_H (see Table 1).

Experiment 2 setting

We used the same setting as described for Experiment 1: two circular arenas, each one with a robot in the centre (fig. 3). While the active robot played the vibration patterns in the action sequence, the passive robot did nothing during the corresponding 30 s. The action sequence that the robots played

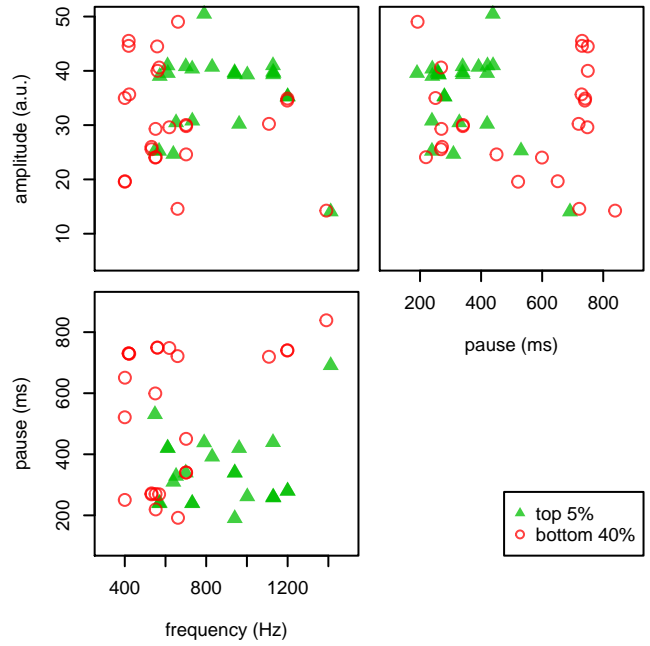


Figure 6: Scatter plot matrix showing the genotypes that obtained a fitness in the top 5%, $e(c) > 34.8$, or in the bottom 45%, $e(c) < -5.8$, of the fitness domain. While there is some clustering in allele values in the best genotypes, there is not a clear separation between the regions occupied by these two extreme sets.

included one genotype from each of the three sets as follows:

$$S = (A, N, V(C_{T_i}), N, A, N, V(C_{L_j}), N, A, N, V(C_{H_k})),$$

with each action lasting for 30 s. This produced three sequences, and we did 10 repeats of each. Each of the 30 experiments used a different set of bees. In this way each group of bees evaluated a genotype from each of the three sets. Such arrangement allowed better comparisons among sets, while minimising the effort of each bee group.

Experiment 2 results

We applied statistical tests to find out genotypes producing similar behaviour to each other. The Welch t-test was applied to the *total bee movement* and *average bee movement* values, since these data show different means and different variances. The F-test was used to compare the *bee acceleration* since the means are zero and variances are different. In both cases we used a significance level of $p = 0.05$. When we examine the result of these statistical tests we see that there are top genotypes that are similar to genotypes with lowest fitness in the last generation and with highest fitness in the first generation. This may be related to the shape of the fitness landscape.

We applied the genotype evaluation function (eq. 1) to the video segment parts (vsps) that correspond to a vibration pattern. We used 58 frames in order to compare the resulting

rank	fitness	fre- quency	pause	amp- litude	remark
1	54.7	1130	260	40	C_T
2	51.3	1130	440	40	C_T
3	48.7	790	440	50	C_T
5	45.0	570	240	25	*
8	43.0	770	420	40	*
9	42.0	500	240	40	*
11	40.7	1030	440	40	C_L
12	40.3	730	240	40	C_H
13	39.3	650	330	30	C_L
41	27.7	680	650	25	C_H
42	26.7	730	390	40	*
58	22.7	1110	170	30	C_L
91	16.3	580	750	30	C_H
112	10.7	730	870	40	*
128	8.7	1300	700	50	*

Table 1: Genotypes used in the post-hoc analysis (C_T , C_L , C_H) and in transferability experiments (*).

values with the fitness values. Figure 7 shows the comparison between values obtained during evolution and the post-hoc experiments. Note that we are plotting the fitness values obtained in each of the three evaluations of a genotype. In the case where a vibration pattern occurred more than once, for comparison we used only the three values that caused such a genotype to be included in one of the C_k sets. We observe that seven out of nine vibration patterns have a similar behaviour in the two experiments. These results confirm that the evolutionary process produces stable and informative results even in a dynamic fitness space mainly due to the variations caused by different bee groups used to produce the fitness values. To further analyse this aspect we compared bee behaviour across segments with no activity.

Table 2 shows the percentage of statistical test results between *no stimuli* segments. In most cases the behaviour is similar but there are a few *vsp*s where the bees are behaving differently, which is unexpected as they are not being subject to any stimuli. It turns out that in most of those cases (in all of them for total and average bee movement) the last *no stimuli* segment of the sequence is involved. This result may be due to bee fatigue (see Mariano et al., 2017).

Transferability of Evolved Genotypes

We were interested in using evolved localised stopping vibration patterns in a different but related task. The goal of this experiment is to analyse if a vibration pattern can be used to aggregate bees in a particular area. We decided to use genotypes with varying fitness values to see if they correlate with aggregation capacity. We selected run number 3, as the genotypes in the first generation had a low fitness, in the middle generations they had a middle fitness value, and




Figure 7: Comparison between fitness evaluations of genotypes obtained in evolution and used in the post-hoc analysis experiments. Chromosomes are sorted by fitness value as used in the ES. Star indicates that the two sets are different as dictated by the Welch t-test ($p < 0.05$). Post-hoc evaluations show a broader distribution but generally confirm fitness values from the evolutionary study.

processed video data	% same	% different with last <i>no stimuli</i> vsp
total bee movement	84.7%	100%
average bee movement	84.5%	100%
bee acceleration	70.7%	61%

Table 2: Results of comparisons between *no stimuli* vsp. Column “% same” indicates the percentage of comparisons with similar results. Last column shows the percentage of comparisons with different results where at least one of the *vsp*s was the last *no stimuli* in the sequence.

the last generation had genotypes with a high fitness value. The genotypes that were selected are the ones marked with a star in column *remark* of Table 1.

Experiment 3 setup

Twelve bees were placed in a long stadium arena that encompasses three robots. Figure 8 shows a picture of the arena superimposed with a mask of the *rois*. The action sequence that the robots played was $S = (N, V)$ with a duration of 180s for each action. We performed at least $n = 12$ repeats for each genotype, each one using a unique bee set.

Experiment 3 results

In order to see if bees are stopping in the active *roi* and if they are aggregating around this region, we analysed plots of NB for each *roi*, for each vibration pattern. We performed a least-squares linear fit of the data of each *roi* and vibration pattern (the data fitted only includes frames where there

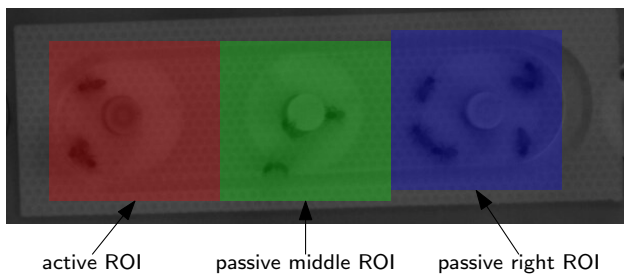


Figure 8: Setup used in the transferability experiments.

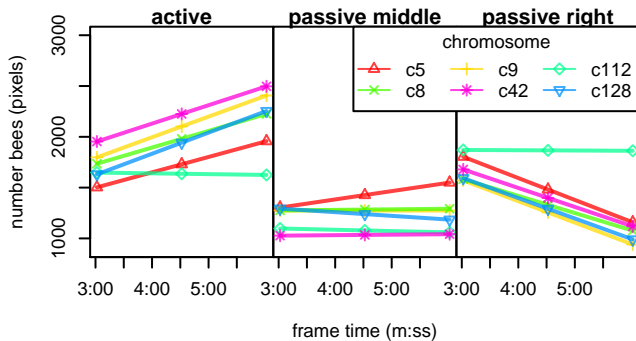


Figure 9: Least-squares fit of number of bees per ROI and per vibration pattern tested, indicating a flow of bees from passive to active regions.

was vibration). Figure 9 shows the result for each ROI. The number of bees generally increases in active ROI at expense of the most distant passive ROI (rightmost) with no substantial change in the closest passive ROI (middle). The evolved vibration pattern successfully attracts bees to the ROI of the vibrating robot.

Discussion and Conclusion

Generally, on-board and on-line evolution of robots that have interactions with living animals is a rather novel field of science. It is posing very high challenges to the established algorithms in the field. As the fitness feedback is derived from living animals, which show very complex and rich nervous and hormonal processing of information, that originates from animal-animal and animal-robot interactions, these systems are essentially black-boxes. The properties of the resulting fitness landscapes are unknown, which situates the problem well for an EC approach although it might be very challenging. One major contribution of our work here is to assess those fitness landscapes in our focal honeybee case and to make first steps to tackle the challenges posed. The fact that living animals are involved stretches the time it takes to evaluate generations in an EC algorithm. Besides the labour-intensity, there are also ethical considerations, that suggest minimizing the number of evaluations to an absolute necessary minimum, which in fact motivates

optimising the efficiency of the experimental procedure and the algorithm used, while simultaneously imposing variability due to needing more than one animal group to obtain the fitness of a single digital genotype.

To the best of our knowledge, this is the first work where evolutionary algorithms were able to evolve robot controllers by using non-human animals as fitness providers. We showed that honeybees guided the evolutionary process of optimising the vibration pattern emitted by robots to stop the bees. We observed that in ten generations the fitness levels out. We observed experimental runs with only little initial improvement of the fitness variables, but then the fitness value stabilised. One reason for these results could be the absence of a gradient in some regions of the fitness landscape, a case in which EC has problems finding the optimum.

An analysis of the evolutionary trajectories in gene space showed that experiments with small improvement in fitness started away from the high fitness region and evolution failed to find this area. This points to a single high fitness region with a plateau around it. One possible solution to escape from a neutral plateau or a local optima is using multi-gene mutation, an approach we should test soon. Both types of difficulty could also be due to the small number of generations with a small population. The dynamic fitness function, resulting from animal fatigue and different group composition, may also play a significant role here. Another effect whose impact requires further investigation is the bee movement evaluation. The current procedure does not distinguish bee linear motion from rotation in place. Therefore, we may be measuring higher movement values than those that really happen. These features are consistent with the differences in fitness observed between the evolution result and the post-hoc evaluation.

In the work presented here, we essentially evolved autonomous robots controllers to produce a stopping signal for living juvenile honeybees. This may be considered a form of unidirectional communication act using a one-word ‘language’, in the sense that the vibration exerts control over juvenile bees, by stopping them. As the evolutionary process unveiled that several parameters of the evolved signal affect the effectiveness of the stop signal – actually leading to a graduated bee motion speed – we probably did not only evolve a single (binary) signal, but also intermediate forms of this signal, bringing a ‘grammar’ aspect into the communication. The transferability experiments indicated that the localised stopping signal can work as an attractor to bees: wandering the bees get ‘caught’ in the ROI of the vibrating robot, which effectively results in the vibration attracting the bees. In the future we plan to go beyond this and include other signals in the evolutionary process, like airflow pulses and temperature pulses that also affect honeybees (Szopek et al., 2013), which can heat up their local environment and which can cool it down with wing beats. This way, we aim for evolving, over time, a set of distinct

signals, ultimately aiming to establish an as-rich-as-possible communication channel between robots and living honeybees.

Acknowledgements

Supported by: EU-ICT project ASSISIBf no. 601074 & centre grant UID/MULTI/04046/2013 from FCT (to BioISI).

References

- Anderson, C. and Ratnieks, F. L. (1999). Worker allocation in insect societies: coordination of nectar foragers and nectar receivers in honey bee (*Apis mellifera*) colonies. *Behav Ecol Sociobiol*, 46(2):73–81.
- Begg, C., Begg, K., Du Toit, J., and Mills, M. (2005). Spatial organization of the honey badger *Mellivora capensis* in the southern kalahari: home-range size and movement patterns. *Journal of Zoology*, 265(1):23–35.
- Bencsik, M., Le Conte, Y., Reyes, M., Pioz, M., Whittaker, D., Crauser, D., Delso, N. S., and Newton, M. I. (2015). Honeybee colony vibrational measurements to highlight the brood cycle. *PLoS one*, 10(11):e0141926.
- Bongard, J. (2008). Behavior chaining-incremental behavior integration for evolutionary robotics. In *ALIFE*, pages 64–71.
- Bonnet, F., Gribovskiy, A., Halloy, J., and Mondada, F. (2018). Closed-loop interactions between a shoal of zebrafish and a group of robotic fish in a circular corridor. *Swarm Intelligence*, pages 1–18.
- Camazine, S., Crailsheim, K., Hrasnigg, N., Robinson, G. E., Leonhard, B., and Kropiunigg, H. (1998). Protein trophallaxis and the regulation of pollen foraging by honey bees (*Apis mellifera* L.). *Apidologie*, 29(1-2):113–126.
- 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.
- Deneubourg, J.-L., Aron, S., Goss, S., and Pasteels, J. M. (1990). The self-organizing exploratory pattern of the argentine ant. *Journal of insect behavior*, 3(2):159–168.
- 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.
- Griparic, K., Haus, T., Miklic, D., and Bogdan, S. (2017). A robotic system for researching social integration in honeybees. In *PLoS ONE* 12(8): e0181977.
- Halloy, J., Sempo, G., Caprari, G., Rivault, C., Asadpour, M., Tâche, F., Said, I., Durier, V., et al. (2007). Social integration of robots into groups of cockroaches to control self-organized choices. *Science*, 318(5853):1155–1158.
- Harvey, I., Husbands, P., Cliff, D., Thompson, A., and Jakobi, N. (1997). Evolutionary robotics: the Sussex approach. *Robotics and autonomous systems*, 20(2-4):205–224.
- Landgraf, T., Oertel, M., Kirbach, A., Menzel, R., and Rojas, R. (2012). Imitation of the honeybee dance communication system by means of a biomimetic robot. In *Biomimetic and Biohybrid Systems*, pages 132–143. Springer.
- Li, W., Gauci, M., and Groß, R. (2013). A coevolutionary approach to learn animal behavior through controlled interaction. In *Procs GECCO*, pages 223–230. ACM.
- Mariano, P., Salem, Z., Mills, R., Zahadat, P., Correia, L., and Schmickl, T. (2017). Animal-guided evolutionary computation in honeybees and robots. In *Procs ECAL*, pages 529–536.
- Michelsen, A., Kirchner, W. H., and Lindauer, M. (1986). Sound and vibrational signals in the dance language of the honeybee, *Apis mellifera*. *Behav Ecol Sociobiol*, 18(3):207–212.
- Nolfi, S. and Floreano, D. (2000). *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. MIT press.
- Ramsey, M., Bencsik, M., and Newton, M. I. (2017). Long-term trends in the honeybee whooping signal revealed by automated detection. *PLoS one*, 12(2):e0171162.
- Reaney, L. T., Sims, R. A., Sims, S. W., Jennions, M. D., and Backwell, P. R. (2008). Experiments with robots explain synchronized courtship in fiddler crabs. *Current Biology*, 18(2):R62.
- Schmickl, T. and Karsai, I. (2016). How regulation based on a common stomach leads to economic optimization of honeybee foraging. *Journal of theoretical biology*, 389:274–286.
- Schmickl, T., Szopek, M., Bodi, M., Hahshold, S., Radspieler, G., Thenius, R., Bogdan, S., Miklic, D., Griparic, K., Haus, T., et al. (2013). ASSISI: Charged hot bees shakin’ in the spotlight. In *Procs. SASO*, pages 259–260.
- Seeley, T. (1994). Honey bee foragers as sensory units of their colonies. *Behav Ecol Sociobiol*, 34(1):51–62.
- Silva, F., Duarte, M., Correia, L., Moura Oliveira, S., and Lyhne Christensen, A. (2016). Open Issues in Evolutionary Robotics. *Evolutionary Computation*, 24(2):205–236.
- Simpson, J. and Cherry, S. M. (1969). Queen confinement, queen piping and swarming in *Apis mellifera* colonies. *Animal Behaviour*, 17, Part 2:271 – 278.
- Stefanec, M., Szopek, M., Schmickl, T., and Mills, R. (2017). Governing the swarm: Controlling a bio-hybrid society of bees robots with computational feedback loops. In *Procs IEEE SSCI*, pages 1–8.
- Swain, D. T., Couzin, I. D., and Leonard, N. E. (2012). Real-time feedback-controlled robotic fish for behavioral experiments with fish schools. *Proceedings of the IEEE*, 100(1):150–163.
- 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.
- Vaughan, R., Sumpster, N., Henderson, J., Frost, A., and Cameron, S. (2000). Experiments in automatic flock control. *Robot Auton Syst*, 31(1):109–117.
- Zhang, B.-T. and Cho (1999). Coevolutionary fitness switching: Learning complex collective behaviors using genetic programming. *Advances in genetic programming*, 3:425–445.