Design choices for adapting bio-hybrid systems with evolutionary computation

Pedro Mariano
BioISI, Faculdade de Ciências, Universidade de Lisboa
plmariano@ciencias.ulisboa.pt

Ziad Salem
Artificial Life Laboratory, Department of Zoology Karl-Franzens-University
ziad.salem@uni-graz.at

Rob Mills
BioISI, Faculdade de Ciências, Universidade de Lisboa
rob.mills@ciencias.ulisboa.pt

Payam Zahadat
Artificial Life Laboratory, Department of Zoology Karl-Franzens-University
payam.zahadat@uni-graz.at

Luís Correia
BioISI, Faculdade de Ciências, Universidade de Lisboa
Luis.Correia@ciencias.ulisboa.pt

Thomas Schmickl
Artificial Life Laboratory, Department of Zoology Karl-Franzens-University
thomas.schmickl@uni-graz.at

ABSTRACT

In this paper we report ongoing work evolving bio-hybrid societies. We develop robots that are integrated into an animal society and accepted as conspecifics. We are using evolutionary algorithms to optimise robot controllers to affect the behaviour of animals. Fitness evaluation is the result of measuring the effect a robot controller has on these animals. Animal habituation and heterogeneous response are two factors that have a major role in this fitness evaluation. We discuss our choices in designing a fitness evaluation procedure and how using animals as fitness function providers impacts this.

ACM Reference format:
Pedro Mariano, Ziad Salem, Rob Mills, Payam Zahadat, Luís Correia, and Thomas Schmickl. 2017. Design choices for adapting bio-hybrid systems with evolutionary computation. In *Proceedings of GECCO '17 Companion, Berlin, Germany, July 15-19, 2017,* 2 pages.
DOI: http://dx.doi.org/10.1145/3067695.3076044

1 INTRODUCTION

Previous approaches to populations of robots interacting with animal populations have in common that they require an *'a priori'* behavioural model of the animals in order to design the robotic devices. In our approach within ASSISIbf [6, 8] we aim to get rid of the need for such an *a priori* model and use computational methods to develop the needed proximate mechanisms of robots in a model-free way during runtime, following ideas that were recently suggested for robotics in a different context (without animal interaction) [2]. As a first step, we here provide the robots with a key signal known from honeybees, the 'queen piping signal' [5], abstracted to a simple series of vibrational pulses intermitted by a significant break of vibration. We identify 2 key parameters that influence the effect of the signal (frequency of vibration & pause length) and use evolutionary computation to 'evolve' a signal with a maximum effect on young honeybees' motion behaviour. We observe the honeybees by automated video analysis and feed back the behaviours as fitness values to inform the optimisation process.

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 our experimental procedure. This article's aim is to explain those pitfalls to allow other scientists to avoid them, as they approach this emerging field of bio-hybrid robotic research. Ultimately we identify the best approach we found so far, in order to provide a first recipe for this novel type of research.

1.1 Related research

The concept of so-called 'bio-hybrid' or 'mixed' societies, involving living organisms and artificial devices, can be traced back to Goumopoulos et al. in 2004 [4] integrating plants with sensors and actuating gadgets, and shortly after Caprari et. al [1] integrating cockroaches and robots (2005). However, no evolutionary experiment had been attempted so far in such societies. It is interesting to contemplate a future co-evolutionary system in which both populations evolve (see [3] for an experiment in a simulated setup where the animals are humans). In our evolving bio-hybrid society only the robot controllers evolve, while the animals do not. Therefore, we consider the animal population to provide fitness information to the evolutionary algorithms developing the robot programs.

2 EVOLUTIONARY PROBLEM

We are evolving a vibration pattern to make bees stop in the active region of interest (ROI). 12 bees are located in a stadium shaped arena with two robots. The robots can act as active (emitting vibration) or passive (generating no stimuli). The chromosome consists of vibration frequency, vibration intensity and pulse period. The evolutionary algorithm (EA) parameters are: population size 5; mutation based on adding Gaussian noise; $\mu + \lambda$ selection; 10 generations. One evaluation trial consists of playing the vibration for 30 s followed by 30 s of airflow. We process a video recording of the first 30 s. From each video frame we compute the number of pixels that differ from: 1) a background image of the experimental arena without bees; 2) a previous frame. These two measures are taken by a ROI, each one centred on a robot. They are a proxy for where the bees are and for how fast they are moving, respectively. The fitness value is the average of processed video values from $N = 3$ evaluation trials. A group of bees is used in at most 30 evaluation trials.

GECCO '17 Companion, Berlin, Germany

DOI: http://dx.doi.org/10.1145/3067695.3076044

3 DESIGN CHOICES

Using animals as fitness function providers poses several challenges when designing an EA.

Animal Habituation. When animals are presented a given stimulus for a long duration, they become habituated and ignore it. We have observed that bees showed a stronger response in initial seconds of the 30 s vibration compared to the last seconds. Since EAs only require a fitness gradient, instead of summing frame-based data, one could instead compute the rate of change in this data.

Heterogeneity of animal response. One challenge of working with animals is their diverse behavioural repertoires. Young animals may exhibit fewer behaviours, 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.

Chromosome Evaluation. To evolve a robot controller we have to take into consideration the possibility of the evolutionary algorithm exploiting unforeseen and idiosyncratic properties of the robot (see [7]). One possibility to mitigate this problem is to evaluate a 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 bees, 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 cater for bee fatigue, i.e., if tired bees were detected in any arena, that arena should be avoided until it is replenished with fresh bees.

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 evaluation trials are for evaluating chromosome c_1, then the second N trials 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 an evaluation trial of chromosome c_1 with the evaluation trial of other chromosomes, to partially mitigate overfitting of a specific bee set to one chromosome. Behavioural heterogeneity can be addressed by increasing the number of evaluation trials, removing outliers, or computing the median.

4 CONCLUSIONS

We have discussed the impacts and problems of using animals as fitness function providers, and our design choices to overcome the problems. Two solutions are increasing and parallelising the number of evaluation trials. The first one comes with a tradeoff of time and animal supply. In the second one, we have to guarantee isolation among arenas. Despite the problems, post hoc analysis of the best chromosomes and the worst chromosomes, of an experiment with a single robot/ROI, have shown that the vibration patterns are capable of stopping the bees and fail to stop them, respectively (see Figure 1). This indicates that our choices of fitness function and evaluation proceeding are good, as we are able to evolve a stopping vibration pattern that affects the entire arena.

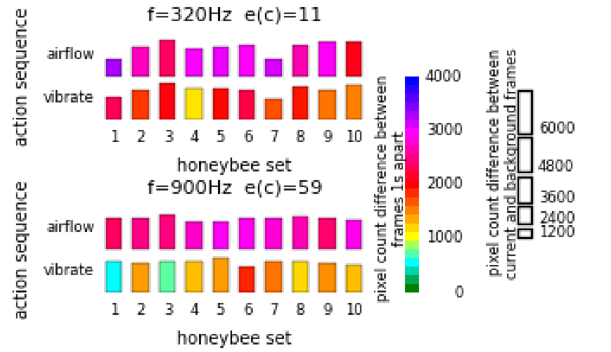

Figure 1: Evaluation metrics from low, $e(c) = 11$, and high, $e(c) = 59$, scoring chromosomes from an evolutionary experiment with a single robot/ROI.

5 FUTURE WORK

Based on these findings we will further explore model-free evolution of more complex animal-robot interaction concerning the signal structure that can be 'shaped' by computational algorithms. This will also involve other signalling modalities beyond pure vibration, e.g. local temperature modulation and airflows. In addition, we will 'evolve' multiple signals that are emitted by each robot in a context-sensitive way, and which will involve multiple cooperating (interacting) robots to generate larger spatial signal patterns. Finally, this research will lead to an evolving adaptive self-organizing bio-hybrid society of autonomous robots and eusocial animals.

ACKNOWLEDGEMENTS

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

REFERENCES

[1] Gilles Caprari, Alexandre Colot, Roland Siegwart, José Halloy, and J-L Deneubourg. 2005. Animal and robot mixed societies: building cooperation between microrobots and cockroaches. *IEEE Robotics & Automation Magazine* 12, 2 (2005), 58–65.

[2] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. 2015. Robots that can adapt like animals. *Nature* 521, 7553 (2015), 503–507.

[3] Pablo Funes, Elizabeth Sklar, Hugues Juillé, and Jordan Pollack. 1998. Animal-animat coevolution: Using the animal population as fitness function. *From Animals to Animats* 5 (1998), 525–533.

[4] Christos Goumopoulos, Eleni Christopoulou, Nikos Drossos, and Achilles Kameas. 2004. *The PLANTS System: Enabling Mixed Societies of Communicating Plants and Artefacts.* Springer, 184–195.

[5] Axel Michelsen, Wolfgang H Kirchner, Bent Bach Andersen, and Martin Lindauer. 1986. The tooting and quacking vibration signals of honeybee queens: a quantitative analysis. *Journal of Comparative Physiology A* 158, 5 (1986), 605–611.

[6] Thomas Schmickl, Stjepan Bogdan, Luís Correia, Serge Kernbach, Francesco Mondada, Michael Bodi, Alexey Gribovskiy, Sibylle Hahshold, Damjan Miklić, Martina Szopek, Ronald Thenius, and José Halloy. 2013. *ASSISI: Mixing Animals with Robots in a Hybrid Society.* Springer, Berlin, Heidelberg, 441–443.

[7] Adrian Thompson. 1996. An evolved circuit, intrinsic in silicon, entwined with physics. In *International Conference on Evolvable Systems.* Springer, 390–405.

[8] Payam Zahadat, Michael Bodi, Ziad Salem, Frank Bonnet, Marcelo Elias de Oliveira, Francesco Mondada, Karlo Griparic, Tomislav Haus, Stjephan Bogdan, Rob Mills, Pedro Mariano, Luís Correia, and Thomas Schmickl. 2014. Social adaptation of robots for modulating self-organisation in animal societies. 2nd FoCAS Workshop on Fundamentals of Collective Systems. (2014).