

VU Research Portal

Combining Conflicting Environmental and Task Requirements in Evolutionary Robotics

Haasdijk, E.W.

published in

Ninth IEEE International Conference on Self-Adaptive and Self-Organizing Systems
2015

DOI (link to publisher)

[10.1109/SASO.2015.21](https://doi.org/10.1109/SASO.2015.21)

document version

Early version, also known as pre-print

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Haasdijk, E. W. (2015). Combining Conflicting Environmental and Task Requirements in Evolutionary Robotics. In *Ninth IEEE International Conference on Self-Adaptive and Self-Organizing Systems* (pp. 131-137). IEEE Press. <https://doi.org/10.1109/SASO.2015.21>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Combining Conflicting Environmental and Task Requirements in Evolutionary Robotics

Evert Haasdijk
Dept. Of Computer Science
VU University Amsterdam
Amsterdam, The Netherlands
Email: e.haasdijk@vu.nl

Abstract—The MONEE framework endows collective adaptive robotic systems with the ability to combine environment- and task-driven selection pressures: it enables distributed online algorithms for learning behaviours that ensure both survival and accomplishment of user-defined tasks. This paper explores the trade-off between these two requirements that evolution must establish when the task is detrimental to survival. To this end, we investigate experiments with populations of 100 simulated robots in a foraging task scenario where successfully collecting resources negatively impacts an individual’s remaining lifetime. We find that the population remains effective at the task of collecting pucks even when the negative impact of collecting a puck is as bad as halving the remaining lifetime. A quantitative analysis of the selection pressures reveals that the task-based selection exerts a higher pressure than the environment.

I. INTRODUCTION

Imagine a collective of robots that is released in an uncharted, possibly changing, environment. The robots have to learn to operate in that environment, of which the particulars are unknown at design time. Thus, the robots have to adapt to circumstances as they find them. Of course, the robot collectives must also serve the purpose of its designers, and so must satisfy their preferences and tasks as well. Typical examples for such scenarios include monitoring, patrolling, surveying, mining or harvesting in remote, inaccessible and possibly hostile environments where human oversight is unfeasible like space, deep mining or undersea [1]. We envision such robot collectives that autonomously adapt through evolution: they evolve controllers that enable them to survive and to perform their tasks.

The environment in which robots operate indirectly circumscribes goals for the population of robots to survive and evolve, but does so without specifying objective functions. Robots must, for instance, move about to spread their genomes, or they must maintain their energy levels by regularly visiting charging stations, but these goals need not be defined explicitly: it is just that robots that display this behaviour get more opportunities to procreate. By virtue of its similarly unbounded nature, biological evolution has resulted in the high levels of adaptability and robustness that we see in natural living organisms. To exploit this creative potential in a system of evolving robots (or robot controllers), we would want to give evolution as much freedom as possible, pushing for open-ended, unbounded adaptivity, unconstrained by user-defined

objective functions.

On the other hand, if the system is to be of any practical relevance, the robots must of course also perform user-defined tasks, pushing for specific, crisply defined task-related objectives.

Evolution has been employed to achieve both of these facets. Artificial Life research abounds with examples of objective-free evolutionary systems since the 1980s [2], [3]. In such experiments, evolution serves as a force for adaptation. Evolutionary robotics research typically employs evolution as a force for optimisation when it focusses on the task-driven aspect [4].

Balancing these two aspects of evolution –environment-driven adaptation and task-driven optimisation– represents a vital step towards implementing our vision of autonomous, functional, responsive and self-sufficient robot collectives. The autonomy that our vision implies prohibits centrally orchestrated evolution, so genomes and performance must be assessed, exchanged and used for selection locally, by the robots themselves.

Earlier work presented the MONEE (Multi-Objective aNd open-Ended Evolution) to solve the problem of combining objective-free and task-driven evolution in a single algorithmic framework [5]. The principal idea behind MONEE is to employ concurrently two selection mechanisms in different roles: environmental selection for open-ended evolution and parent (or mate) selection for task-driven adaptation, both operating solely at local level. The results reported there show that the strategy of adding explicit task-based selection to an environment-driven evolutionary system yields a system where robots evolve behaviour that allows them to procreate in the environment as well as perform their tasks. These results were based in a scenario where the task (resource gathering) did not interfere with the demands of the environment. Environmentally optimal behaviour required moving around the arena as much as possible, and this aligned with optimal task behaviour, which involved seeking and collecting pucks strewn throughout the environment: movement was commonly beneficial behaviour without need for compensation.

This paper investigates the effects of a scenario where the task requirements oppose those of the environment. In particular, cases where executing the task is detrimental to a robots life expectancy, e.g., because it implies a physical risk

or simply drains batteries rapidly. The environment and task are essentially the same as in [5], but in this case, picking up a puck incurs a penalty and the robot's remaining lifetime is reduced by some percentage – the pucks are poisonous. This causes a complex interaction between the requirements posed by task and environment. Disregarding the task, the robots would avoid the pucks to maximise their lifetime and so have more opportunities to spread their genome. However, robots that perform the task well are more likely to be selected as parents once they have spread their genomes.

We conduct a series of simulation experiments with populations of 100 robots and varying poison levels to investigate what balance evolution strikes between these conflicting demands. In the following, we show that the explicit task-based parent selection outweighs the implicit environmental survivor selection until the penalty of performing the task becomes very high.

II. RELATED WORK

Evolutionary Robotics has been widely studied since the early 1990s as a tool to design the morphology and control of robots through evolutionary algorithms [4]. Initially, research focussed on individual robots, but since then substantial effort has been directed at evolution in larger numbers of interacting autonomous robots in swarms [6], research projects include for instance the Swarmanoid project [7]) or modular robots (e.g. M-tran [8]). Bongard [9] provides an overview of this vein of evolutionary robotics research. Evolutionary robotics research encompasses evolutionary algorithms to develop straightforward tasks such as obstacle avoidance for differential drive robots [4] to Bongard et al's artificial ontogeny [10] that develops morphology and control in concert and from evolving diverse behaviours [11] to self-modelling [12]. What almost all these contributions to the field have in common is that the evolutionary process is employed to optimise robots to achieve some fixed user-defined objective at design time. They employ the classical evaluate-select loop of evolutionary computing and in that sense differ little from non-robotic evolutionary algorithms.

Watson et al. [13] coined the phrase 'embodied evolution' in a depart from this paradigm to one of distributed on-line evolution that enables adaptivity at run time. With embodied evolution, the robots autonomously find mates and exchange genetic material without central oversight: evolution occurs continuously, asynchronously and at run time through robot interactions rather than at design time by virtue of a central control loop. Selection is based on local assessment of task performance, so it shares the task-driven focus of more mainstream evolutionary robotics. Embodied evolution -on-line evolution in general- implies some environmental selection: robots that do not meet the requirements of the environment and end up stuck or with empty batteries effectively remove themselves from the gene pool. The implications of this environmental pressure are not considered, however, nor are they separately researched. In many cases, embodied evolution

implementations employ task-related (virtual) energy to determine parent as well as survivor selection [13], [14]. In these cases, there is no clear distinction between selection based on task performance and based on environmental aptitude and task and environment are necessarily aligned.

Objective-free evolution as well as self-replication have been studied in Artificial Life since Rasmussen's [15] and Ray's [16] work. Such research primarily investigates evolutionary dynamics in the absence of tasks, but as a result of implicit or environmental criteria that impact the ability to spread genomes through the population. Such open-ended approaches have gained interest from the evolutionary robotics community, for instance in Bianco and Nolfi's experiments with self-assembling organisms [17], Schwarzer et al's work on artificial sexuality [18] and more recently in the MEDEA algorithm [19].

Bredecche et al. describe MEDEA [19], an open-ended evolutionary algorithm where autonomous robots move around an arena while continually broadcasting their genome over a short range. Meanwhile, they also receive genomes from other robots that come in communication range. When a robot's lifetime expires, it randomly selects one of the received genomes, modifies that using mutation and starts a new life of broadcasting this new genome. This set-up promotes, with only environmental selection, robot movement through the environment: genomes that cause the robot to move around a lot are spread at a much higher rate than genomes that cause their host to stand still.

In the Avida Artificial Life system [20], organisms – 'Avidians', virtual machine code programs– face a combination selection pressures. One derives from the environment and considers only efficient replication. The other source of pressure is a task: Avidians that correctly solve some computational problem are rewarded through an increased rate of execution. Task-based selection is implicit: faster execution implies more opportunities to procreate.

MONEE extends objective-free approaches with a system where an individual can accumulate credits through task performance – the better a robot performs a task, the more credits it earns [5]. When an individual puts its genome forward as a potential parent, it also passes information on its earnings as an indication of its worth. The genomes with the highest associated credits are then selected to produce new offspring (inspired by [18], but there an individual's capital was fixed and did not reflect proficiency at any task).

The MONEE scheme is reminiscent of parental investment, which has been investigated in ALife settings, including experiments with robots [18], [21], [22]. In artificial life parental investment is often used to give the offspring a starting value of (virtual) energy [23]–[26] and a parent's energy level is often linked to task performance (e.g., agents tasked with eating grass to gather energy [25]). The MONEE scheme differs subtly but crucially from such parental investment schemes: a parent does not actually invest when impregnating an egg because the credits aren't transferred but copied; there is no cost involved.

III. MONEE: MULTI-OBJECTIVE & OPEN-ENDED EVOLUTION

A. Environmental Selection

The robot controller lifecycle in our experiments consists of two phases: life and rebirth. The robot controllers have a limited, fixed, lifetime during which they perform their actions; moving about, foraging, et cetera (this lifetime may be reduced by picking up pucks as described below). When their lifetime ends, they enter a rebirth phase and become ‘eggs’: stationary receptacles for genomes that are transmitted by passing live robots.

The rebirth phase also lasts a fixed amount of time, and once this has passed, the egg selects parents from the received genomes to create a new controller. The robot then reverts to the ‘life’ role with this new controller. Thus, robot controllers can procreate by transmitting their genome to eggs, and the more eggs a robot inseminates, the more chances it has for procreation. Because the transmission of genomes is continuous and at close range (e.g. through infrared), the more a robot moves about the arena, the better its chances of producing offspring.

This defines the environmental part of selection that promotes movement and is based on the experiments with *meDEA* by Bredèche et al. [19]. This aspect of the set-up is open-ended in the sense that it is objective-free: there is no calculated performance measure that defines the chances of being selected as parent, there is no task. Only the interaction between environment and robot behaviour dictates which individuals may or may not become parents.

B. Task-based Selection

To add task-driven parent selection to this basic evolutionary process, the robots can, during their lifetime, amass credits by performing tasks. For instance, a robot could get one credit for every piece of ore it collects, one for successfully solving some puzzle, and so on. When a robot passes its genome to an egg, it passes the current credit count along with it and the egg uses that information to select parents when it revives. The credits relate task performance to reproductive success: besides the open-ended goal of ‘merely’ transmitting genomes to eggs, robots must also become proficient at the defined tasks for these genomes to be selected. The more proficient a robot is at a task, the higher its chances of procreating, but only when it also successfully negotiates the environment.

Thus, we define two stages of selection: the first is implicit and derives from the definition of the environment and its rules for genome transfer. In the set-up of our experiments, it promotes movement, just as in *meDEA*. The second selection stage explicitly considers the awards amassed by the individual when selecting from genomes that each egg collected.

IV. EXPERIMENTAL SET-UP

A. Environment and Task

We implemented the MONEE algorithm in a simple 2D simulator called *RoboRobo* [27], simulating 100 e-puck robots

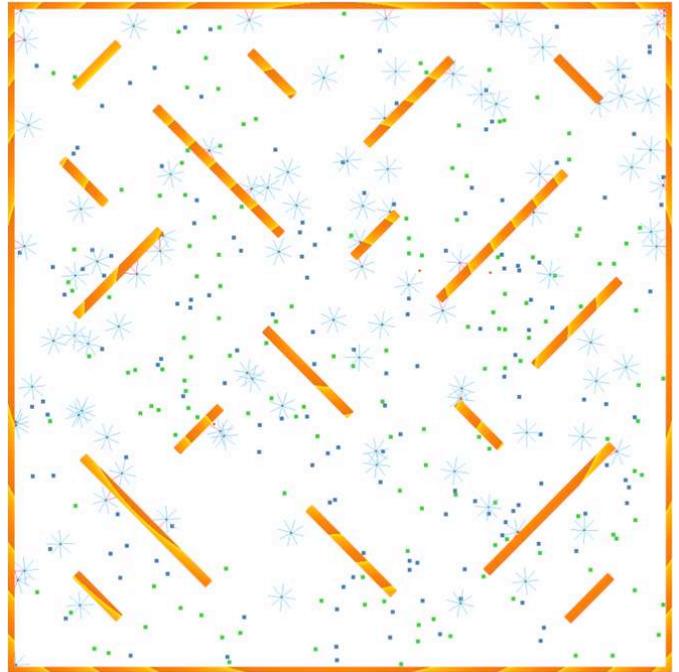


Fig. 1. Experiment screenshot. Robots are shown as small circles with sensor beams indicated. Pucks are shown as small green squares (the blue squares show a second puck type that is disregarded in the experiments in this paper). The shaded orange rectangles indicate arena walls and obstacles.

in an environment that contains obstacles and pucks.¹ The sides of the square arena are roughly 330 robot body lengths long (1024 pixels in the simulator), and it contains a number of obstacles (see Fig. 1) and pucks. The pucks are spread throughout the arena, and they are immediately replaced in a random location when picked up. The robots move around the arena, spreading their genome as they encounter eggs and dying when their allotted time has passed.

Robots collect pucks simply by driving over them; the more pucks they gather, the more likely their genome is to be selected once an egg they impregnated revives. To detect pucks, the robots have 8 special sensors, laid out in the same manner as the standard e-puck infrared sensors: 6 face forward, 2 face to the rear. Each robot is controlled by a single-layer feed forward neural network which controls its left and right wheels. The inputs for the neural network are the robot’s puck and obstacle sensors as well as two bias nodes. The robot’s genome directly encodes the neural network’s weights as an array of reals.

As mentioned, the robots alternate between periods of active puck gathering (life phase) and motionless genome reception (egg phase). The egg phase lasts 200 time steps, the life phase is initialised at 2,000 time steps, but to prevent synchronised cycles among the robots, we add a small random number to each robot’s fixed lifetime. This desynchronises switching between life and rebirth even though our runs start with all robots in sync at the first time-step of their lifetime.

¹Code for the experiments and analysis scripts is available from <https://github.com/ci-group/monee.git>.

At the end of the egg phase, a parent is selected from the received genomes using binary selection on the basis of the number of pucks collected (i.e., on their task performance). Offspring is created by and mutating the parent weights using gaussian perturbation with a single, fixed mutation step size $\sigma = 1$. This single-parent, mutation-only scheme is common in evolution strategies that are known to perform well on problems with continuous-valued genomes [28].

B. Poisonous Pucks

To juxtapose the task and environmental requirements, collecting a puck incurs a penalty: the robot’s remaining lifetime is reduced by a set percentage – the poison level. Suppose, for instance, that some robot has a remaining lifetime of 1,000 time steps as it picks up a puck in an experiment where the poison level is set to 10%. It’s remaining lifetime is then immediately reduced to 900 time steps. If this robot were to immediately pick up another puck, its remaining lifetime would again be reduced by 10%, this time to 810 time steps. In absolute numbers, this penalises ‘young’ individuals more than it does ‘old’ ones, but the robots have no concept of time or age and so cannot take this into account, for instance, to start gathering pucks later in life. We run a series of experiments with poison levels fixed at 0, 2, 5, 10, 20 or 50%.

As control, we also run experiments where no task is defined and parent selection within an egg is random, disregarding the number of pucks collected. We run 32 repetitions of each experiment.

C. Quantifying Selection Pressure

To facilitate a quantitative comparison of the selection pressure exerted by the environment and by the task, we use Kendall’s τ coefficient [29] to quantify the correlation between behaviour and number of offspring in a population. High values for τ indicate a strong correlation and therefore a high selection pressure, allowing us to quantitatively compare the selection pressure between different scenarios.

Because the robots change controllers asynchronously, the definition of a population is not entirely straightforward. In this case, we define a population as all the controllers that ran to completion within a certain timeframe (intervals of 5000 time steps).

When parent selection is based on the number of pucks collected, we calculate τ for a population to determine the correlation between the number of pucks collected to the number of children for each individual. The control experiments require another measure of relevant behaviour as the number of collected pucks is explicitly disregarded in these experiments.

The rules for genome exchange create a pressure towards movement in this case, so the most relevant behavioural measure is the total distance an individual has covered during its life. Again, we measure the correlation, this time between distance covered and offspring.

V. RESULTS

To assess the take-up of the tasks, we need to establish whether the robots learn to gather pucks when the task is

in force. Figure 2 shows the number of pucks collected over time for different poison levels with and without the task. The left-hand panel shows a clear trend towards collecting increasing numbers of pucks as evolution progresses. The grey line shows the number of pucks collected when the pucks are not poisonous, the blue lines show the same for increasing poison levels.

Figure 2b shows that the robots collect far fewer pucks in the control experiments. Even when the pucks are not poisonous (grey line), the number of pucks collected is much lower than for the highest poison level with the task. For increasing poison levels the robots learn to avoid the pucks so as to maximise their lifetime (blue lines).

Clearly, even when collecting pucks has a substantial impact on lifetime, the evolutionary balance still favours the task. The number of pucks collected does decrease as the poison level rises. This is at least in part due to the fact that the reduction in lifespan that results from collecting poisonous pucks. The shorter lifespan causes a larger amount of time to be spent in egg state, when no collecting takes place.

To correct for this phenomenon, we also consider the number of pucks collected normalised for lifetime by calculating the ratio of pucks collected to lifetime: an individual that lived for T time steps and collected N pucks has a ratio of $\frac{N}{T}$; the number of pucks collected per tick of activity. Figure 3 plots this ratio over time. This shows that the robots, in fact, barely reduce their puck-collecting efforts - the rate at which robots pick up pucks while they are active develops almost identically, or at least reaches identical levels, for poison levels up to 20% as it does when the pucks are not poisonous at all. Only when the poison level reaches 50%, i.e., when picking up a puck *halves* the remaining lifespan, do we see an appreciable effect.

It appears that the evolutionary equilibrium between task and environment is very biased to favour the task, at least for the range of poison levels we considered. We were surprised by this finding: the robots spread their genomes at a much lower rate when their lifetime is shortened, and if an individual’s genome isn’t collected by eggs, what would the benefit of a higher puck count be?

To understand why the task is so much more important for evolution, we compare the pressure exerted by the two selection mechanisms. To this end, we quantify the selection pressure for as described in Section IV-C. We divide the experiment into slices of 5,000 ticks and then consider the robots that complete their lifetime during each interval as a population where we quantify selection pressure in terms of distance covered or pucks collected as described above.

Figure 4 shows the results of this analysis. The trend for the experiments with the task of an initially low selection pressure that peaks and then levels off once the required behaviour is well established in the population is similar to that reported in [5]. The selection pressure in the control experiment (red line) that derives from the environment’s push for movement hovers around $\tau = 0.1$ – if there is a similar trend, it is much less pronounced.

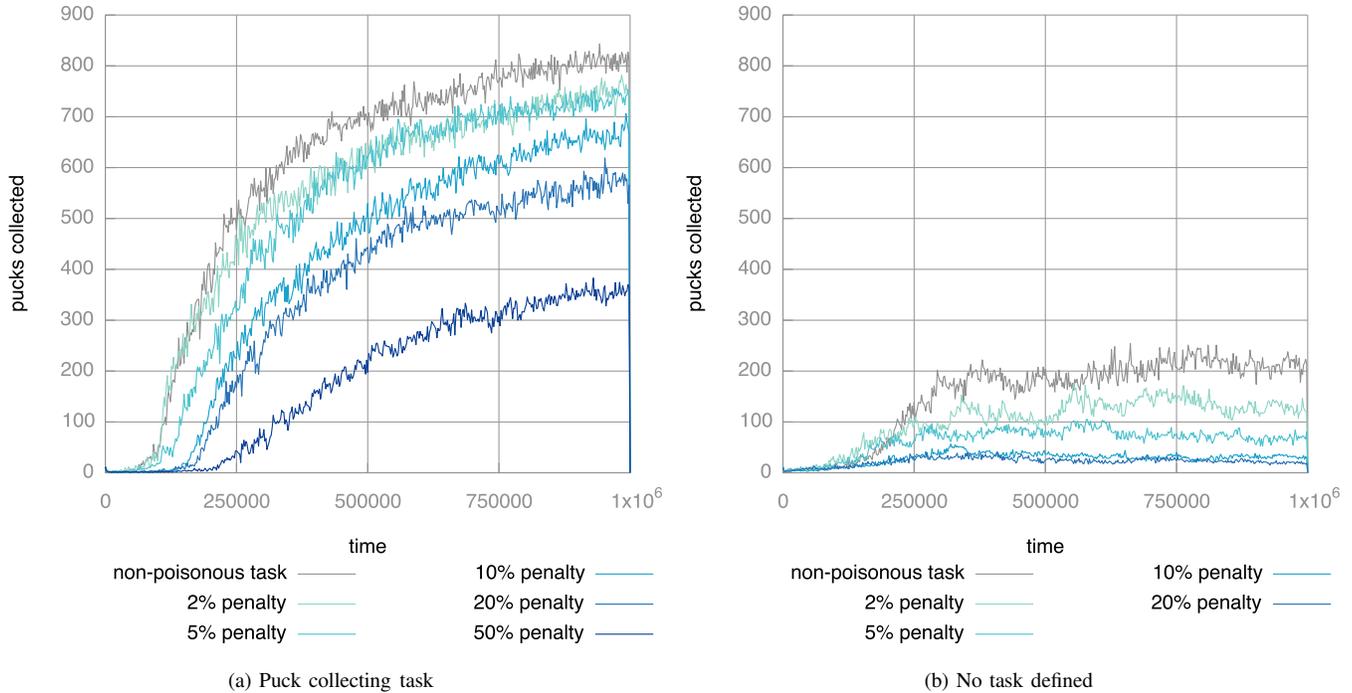


Fig. 2. The median ($N=32$) number of pucks collected by the population per 1,000 time-steps for poison levels ranging from a 0 (grey) to 50 (dark blue) percent lifetime penalty. The left-hand graph shows results when parent selection is based on the number of pucks collected. The right-hand panel shows the number of pucks collected when parent selection is random (i.e., when there is only environmental selection pressure). When the task is in effect, the number of pucks collected does not drop substantially even though collecting pucks impacts the robots’ lifetime. When parent selection is random, far fewer pucks are collected, particularly for higher poison levels.

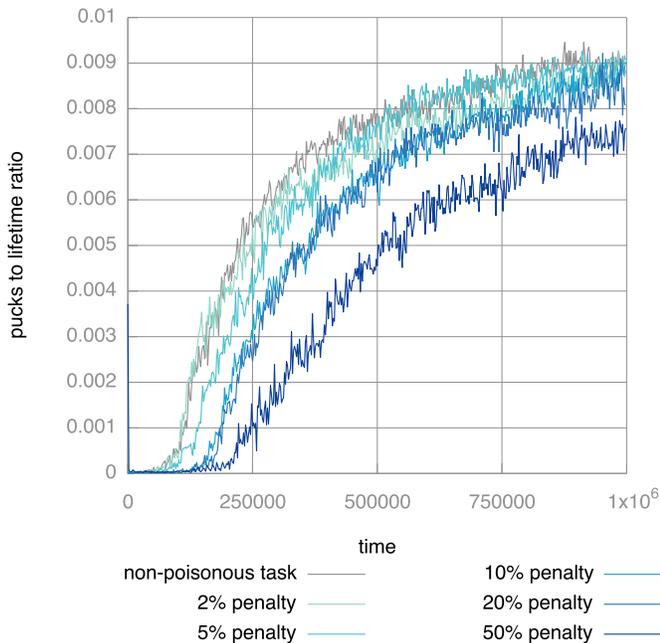


Fig. 3. The median ($N=32$) number of pucks collected, normalised for lifetime, per 1,000 time-steps for poison levels ranging from a 0 (grey) to 50 (dark blue) percent lifetime penalty. This highlights the fact that the poison level has only limited influence on the balance between survival and task.

It seems that the disparity in selection pressure from environment and task results in the bias towards behaviour that is effective for the task, but detrimental in terms of survival. At a poison level of 50% the magnitude of task-based selection pressure is very close to the selection pressure exerted by the environment. Even such a small difference causes substantial bias towards task performance: the decrease in pucks to lifetime ratio between 20 and 50% is clear in Figure 3, and it contrasts with the almost overlapping curves for lower poison levels, but the drop is not dramatic.

VI. CONCLUSION

This paper investigated the problem of combining task-directed optimisation and environment-driven adaptation in situations where task and environment conflict. Specifically, the robots in our experiments were tasked with collecting pucks while this very act reduced their lifespan.

We found that, even when the penalty of collecting a puck amounts to halving an individual’s lifetime, the populations collect substantial amounts of pucks, many more than they do in the control experiments where collecting pucks has no adverse effect nor any benefit. The obvious, if somewhat surprising, conclusion is that the benefit of collecting pucks an increasing one’s chances in the second stage of selection which considers task performance outweighs the benefit of longevity when pucks are avoided.

We quantitatively analysed the selection pressure due to

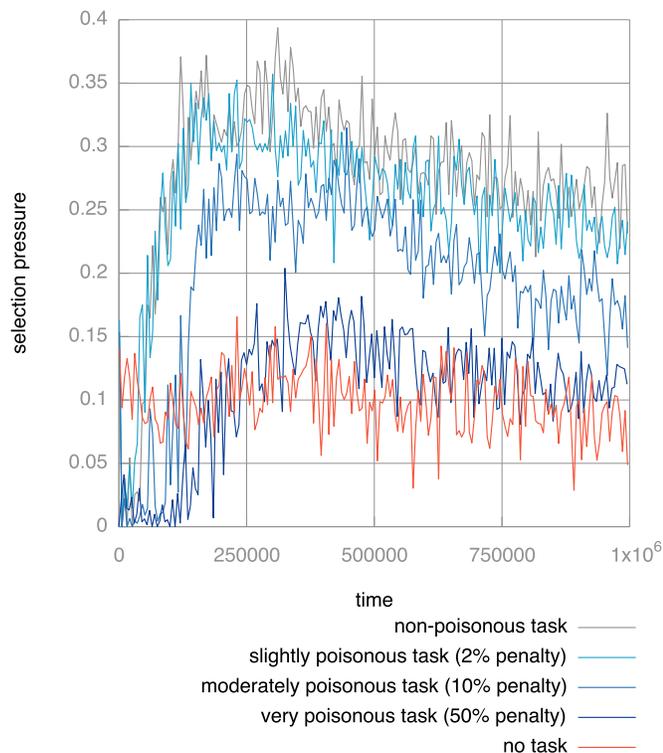


Fig. 4. Quantitative analysis of selection pressure over time. Selection pressure is quantified as the correlation between task performance (when there is no task defined: distance covered during the individual's lifetime) and number of offspring per individual. The selection pressure that derives from the environment (red) is clearly lower than that from task-based parent selection for poison levels up to 20%. For a 50% poison level, the selection pressure from puck collection is still, albeit slightly, higher than the environmental selection pressure.

environment-driven survivor selection and that due to task-based parent selection to better understand the cause of this proclivity for task-based behaviour. The analyses showed that explicitly selecting for task performance (using binary tournament in these experiments) yields a selection pressure that is substantially higher than the selection pressure implied by the environment. The task-based selection pressure does decrease for higher poison levels, but even the small difference in the magnitude of selection pressure we see for a very high poison level (halving the remaining lifetime for every puck collected, or a 50% poison level) is enough to warrant substantial task performance.

This leads to the conclusion that the effect of combining selection mechanisms depends strongly on the relative strengths of the selection pressure these mechanisms exert. Even small differences in the magnitude of selection pressures may cause one selection mechanism to outweigh another when they are combined.

This finding may have implications beyond the current case of environment-driven and task-based selection to other combinations of selection, for instance when selection is partly interactive or in memetic systems where social learning

and evolution interleave. We will investigate the effects of juxtaposed requirements further, for instance also considering conflicting tasks (where selection pressure should be of comparable magnitude).

ACKNOWLEDGMENT

The description of the vision, system and experiments in this paper is based on earlier work (in particular, viz. [5])

REFERENCES

- [1] J. G. Bellingham and K. Rajan, "Robotics in remote and hostile environments." *Science*, vol. 318, no. 5853, pp. 1098–1102, 2007.
- [2] C. G. Langton, *Artificial Life: Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 1989.
- [3] C. Langton, Ed., *Artificial Life: an Overview*. MIT Press, Cambridge, MA, 1995.
- [4] S. Nolfi and D. Floreano, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press, Cambridge, MA, 2000.
- [5] E. Haasdijk, N. Bredeche, and A. E. Eiben, "Combining environment-driven adaptation and task-driven optimisation in evolutionary robotics," *PLoS ONE*, vol. 9, no. 6, p. e98466, 2014. [Online]. Available: <http://www.plosone.org/article/info%3Adoi%2F10.1371%2Fjournal.pone.0098466>
- [6] V. Trianni, *Evolutionary Swarm Robotics – Evolving Self-Organising Behaviours in Groups of Autonomous Robots*, ser. Studies in Computational Intelligence. Berlin–Heidelberg–New York: Springer–Verlag, 2008, vol. 108.
- [7] M. Dorigo, D. Floreano, L. M. Gambardella, F. Mondada, S. Nolfi, T. Baaboura, M. Birattari, M. Bonani, M. Brambilla, A. Brutschy, D. Burnier, A. Campo, A. Christensen, A. Decugnière, G. Di Caro, F. Ducatelle, E. Ferrante, A. Förster, J. Guzzi, V. Longchamp, S. Magnenat, J. Martinez Gonzales, N. Mathews, M. Montes de Oca, R. O'Grady, C. Pinciroli, G. Pini, P. Rétonnaz, J. Roberts, V. Sperati, T. Stirling, A. Stranieri, T. Stützle, V. Trianni, E. Tuci, A. E. Turgut, and F. Vaussard, "Swarmanoid: a novel concept for the study of heterogeneous robotic swarms," *IEEE J. Robot. Autom.*, p. in press, 2012.
- [8] H. Kurokawa, E. Yoshida, K. Tomita, A. Kamimura, S. Murata, and S. Kokaji, "Self-reconfigurable m-tran structures and walker generation," *Robotics and Autonomous Systems*, vol. 54, no. 2, pp. 142–149, 2006.
- [9] J. Bongard, "Evolutionary robotics," *Communications of the ACM*, vol. 56, no. 8, pp. 74–85, 2013.
- [10] J. C. Bongard and R. Pfeifer, "Evolving complete agents using artificial ontogeny," in *Morpho-functional Machines: The New Species*, F. Hara and R. Pfeifer, Eds. Springer Japan, 2003, pp. 237–258. [Online]. Available: http://dx.doi.org/10.1007/978-4-431-67869-4_12
- [11] J. Lehman and K. O. Stanley, "Abandoning objectives: Evolution through the search for novelty alone," *Evol. Comput.*, vol. 19, no. 2, pp. 189–223, Jun. 2011. [Online]. Available: http://dx.doi.org/10.1162/EVCO_a_00025
- [12] J. Bongard, V. Zykov, and H. Lipson, "Resilient Machines Through Continuous Self-Modeling," *Science*, vol. 314, no. 5802, pp. 1118–1121, 2006. [Online]. Available: <http://www.sciencemag.org/content/314/5802/1118.abstract>
- [13] R. A. Watson, S. G. Ficici, and J. B. Pollack, "Embodied evolution: Distributing an evolutionary algorithm in a population of robots," *Robotics and Autonomous Systems*, vol. 39, no. 1, pp. 1–18, Apr. 2002. [Online]. Available: <http://eprints.ecs.soton.ac.uk/10620/>
- [14] S. Wischmann, K. Stamm, and F. Wörgötter, "Embodied evolution and learning: The neglected timing of maturation," in *Advances in Artificial Life: 9th European Conference on Artificial Life*, ser. Lecture Notes in Artificial Intelligence, F. Almeida e Costa, Ed. Lisbon, Portugal: Springer-Verlag, September 10–14 2007, vol. 4648, pp. 284–293.
- [15] S. Rasmussen, C. Knudsen, R. Feldberg, and M. Hindsholm, "The coreworld: Emergence and evolution of cooperative structures in a computational chemistry," *Physica D: Nonlinear Phenomena*, vol. 42, no. 1, pp. 111–134, 1990.
- [16] T. S. Ray, "Is it alive or is it GA," in *Fourth International Conference on Genetic Algorithms*, R. K. Belew and L. B. Booker, Eds., 1991, pp. 527–534.

- [17] R. Bianco and S. Nolfi, "Toward open-ended evolutionary robotics: evolving elementary robotic units able to self-assemble and self-reproduce," *Connection Science*, vol. 4, pp. 227–248, 2004.
- [18] C. Schwarzer, C. Höslér, and N. Michiels, "Artificial sexuality and reproduction of robot organisms," in *Symbiotic Multi-Robot Organisms: Reliability, Adaptability, Evolution*, P. Levi and S. Kernbach, Eds. Berlin–Heidelberg–New York: Springer–Verlag, May 2010, pp. 384–403. [Online]. Available: <http://www.springer.com/engineering/mathematical/book/978-3-642-11691-9>
- [19] N. Bredeche, J.-M. Montanier, W. Liu, and A. F. Winfield, "Environment-driven distributed evolutionary adaptation in a population of autonomous robotic agents," *Mathematical and Computer Modelling of Dynamical Systems*, vol. 18, no. 1, pp. 101–129, 2012.
- [20] C. Adami and C. T. Brown, "Evolutionary learning in the 2d artificial life system avida," *Artificial Life IV*, pp. 377–381, 1994.
- [21] S. Mascaro, K. Korb, and A. Nicholson, "An alife investigation on the origins of dimorphic parental investments," in *Advances in Natural Computation, Proceedings of the Australian Conference on Artificial Life (ACAL 2005)*, H. A. Abbass, T. Bossomaier, and J. Wiles, Eds., vol. 3, 2005, pp. 171–185.
- [22] J. Ventrella, "Genepool: Exploring the interaction between natural selection and sexual selection," *Artificial Life Models in Software*, pp. 81–96, 2005.
- [23] F. Menczer and R. Belew, "Latent energy environments," in *Santa Fe Institute Studies In The Sciences Of Complexity-Proceedings Volume-*, vol. 26, 1996, pp. 191–210.
- [24] F. Menczer, W. Willuhn, and R. Belew, "An endogenous fitness paradigm for adaptive information agents," in *CIKM Workshop on Intelligent Information Agents*. Citeseer, 1994.
- [25] M. Burtsev, V. Red'ko, and R. Gusarev, "Model of evolutionary emergence of purposeful adaptive behavior. the role of motivation," in *ECAL*, ser. Lecture Notes in Computer Science, J. Kelemen and P. Sosík, Eds., vol. 2159. Springer, 2001, pp. 413–416.
- [26] M. Scheutz and P. Schermerhorn, "Predicting population dynamics and evolutionary trajectories based on performance evaluations in alife simulations," in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2005)*, H.-G. Beyer and U.-M. O'Reilly, Eds., ACM. ACM, 2005, pp. 35–42.
- [27] N. Bredeche, J.-M. Montanier, B. Weel, and E. Haasdijk, "Roborobo! a fast robot simulator for swarm and collective robotics," *CoRR*, vol. abs/1304.2888, 2013.
- [28] H.-G. Beyer and H.-P. Schwefel, "Evolution strategies – A comprehensive introduction," *Natural Computing*, vol. 1, pp. 3–52, 2002.
- [29] M. G. Kendall, "A new measure of rank correlation," *Biometrika*, vol. 30, no. 1-2, pp. 81–93, 1938. [Online]. Available: <http://biomet.oxfordjournals.org/content/30/1-2/81.short>