Macroscopic analysis of robot foraging behaviour

Michel van Dartel Eric Postma Jaap van den Herik
Guido de Croon

IKAT/Universiteit Maastricht

P.O. Box 616

6200 MD Maastricht

Tel: +31 (0)43 38 83477

Fax: +31 (0)43 38 84897
{mf.vandartel, postma, herik, g.decroon}@cs.unimaas.nl

Abstract

Microscopic analysis is a standard approach in the study of robot behaviour. Typically, the approach comprises the analysis of a single (or sometimes a few) robot-environment system(s) to reveal specific properties of robot behaviour. In contrast to microscopic analysis, macroscopic analysis focuses on averaged properties of systems. The advantage is that such a property is easier to generalise so that it can be established to what extent the property is universal. This paper investigates whether a macroscopic analysis can reveal a universal property of adaptive behaviour in a robot model of foraging behaviour. Our analysis reveals that the step lengths of the most successful robots are distributed according to a Lévy-flight distribution. From studies on a variety of natural species, it is known that such a distribution constitutes a universal property of foraging behaviour. Thereafter we discuss an example of how macroscopic analysis can be applied to existing research in evolutionary robotics, and relate the macroscopic and microscopic analyses of foraging behaviour to the framework of scientific research described by Cohen (1995). We conclude that macroscopic analysis may predict universal properties of adaptive behaviour and that it may complement microscopic analysis in the study of adaptive behaviour.

1 Introduction

In order to understand adaptive behaviour, Beer (2003) suggested that one should start by investigating minimal models, as is a common approach in physics. In recent years, in-depth analysis of such models revealed many new insights into the processes underlying cognition (see, e.g. Beer, 1995, 2000, 2003, Nolfi, 1998, 2002, Slocum, Downey, and Beer, 2000, van Dartel, Sprinkhuizen-Kuyper, Postma, and van den Herik, in press, de Croon, Nolfi, and Postma, in press). So far, analysis of single robots has been performed at a microscopic level, in which the focus is on the successful behaviour of individual robots. Although such microscopic analyses have led to explanatory insights and testable predictions for individual robots in specific environments, the insights and predictions are difficult to generalise (Guillot and Meyer, 2001). For example, in a study by Nolfi (2002), a microscopic analysis of the sensorimotor behaviour of a few successful robots on a self-localisation task revealed that success on that specific task depended on the ability to integrate information over different time scales. Moreover, microscopic analyses reveal properties of robot behaviour on various tasks. At the same time the microscopic approach raises questions about the generality of the properties revealed. In the case of Nolfi's study, it remains to be established how general the strategy of integrating information over different time scales is in relation to self-localisation. We believe that for adequate handling of such issues of generality another approach is needed.

In this paper we demonstrate that *macro*scopic analysis might be such an approach. Applied to robotics, macroscopic analysis analyses the average properties of many robots in similar environments. So far, macroscopic analysis has only been applied to robotics in the context of multi-robot systems (see, e.g. Lerman, Galstyan, Martinoli, and Ijspeert, 2001, Lerman and Galstyan, 2003, Jones and Mataric, 2003). In order to improve the generality of findings in single-robot systems, macroscopic analysis has to be conducted over many instances of the same robot on the same task.

For two reasons, macroscopic analysis might be adequate to handle issues of generality in evolutionary robotics. The first reason is that macroscopic analysis ignores individual differences by analysing the average properties of many robots in similar environments. Therefore, macroscopic analysis may be more suitable to reveal universal properties, i.e. properties characteristic for a class of robot-environment systems. The second reason is that the application of macroscopic analysis in statistical physics has revealed many universal properties; for instance on DNA sequences, heartbeat rates, and weather variations (Havlin et al., 1999, Stanley et al., 2001, Beekman, Sumpter, and Ratnieks, 2001).

In summary, we propose that macroscopic analysis of single-robot systems might improve the generalisability of contemporary research in robotics. To evaluate the effectiveness of macroscopic analysis in the context of single robot systems we formulated a straightforward research question: *Can macroscopic analysis reveal a universal property of adaptive behaviour in a robot model?*

To answer this research question we employ the methodology of evolutionary robotics (Nolfi and Floreano, 2000: 18), because "the possibility of evolving robots that are free to select their way to solve a task by interacting with their environment may help us to understand how natural organisms produce adaptive behaviour". In our research, we will optimise foraging behaviour of neural-network-controlled robots using an evolutionary-computation algorithm. Then, we will perform a macroscopic analysis on the foraging behaviour of the optimised robots. We expect that, by averaging over many behavioural interactions, macroscopic analysis will reveal properties that would remain obscured when only microscopic analysis was applied. Hence, we investigate whether macroscopic analysis will predict a universal property of foraging behaviour that can be confirmed by microscopic analysis. Additionally, we will perform a microscopic analysis to examine whether a property revealed by the macroscopic analysis can be recognised in the microscopic behaviour of the robots.

The outline of the remainder of the paper is as follows. In section 2, the foraging experiment is described. Section 3 presents the experiments and the results. In section 4 a macroscopic analysis of the robot's foraging behaviour is conducted and a provisional conclusion is provided. Section 5 contains the microscopic analysis and gives another provisional conclusion. In section 6, the macroscopic and microscopic analyses are discussed. Finally, our conclusions are presented in section 7.

2 The foraging experiment

The foraging experiment is described in terms of the environment (subsection 2.1) and the robot (subsection 2.2).

2.1 The environment

The environment is defined as an $L \times L$ square lattice with periodic boundary conditions (i.e. the environment is defined on a torus) containing n food elements. Randomly distributed dots over the L^2 locations represent the food elements. A robot will collect food by walking over the food elements. Whenever a food element is collected, it is removed from the environment and replaced by a new one at a random location. In this way, the number of food elements remains constant throughout the experiment. Figure 1 is an illustration of the environment with randomly distributed food elements (dots), and in the environment a robot (circle).

Insert figure 1 about here.

2.2 The robot

The robot performing the foraging task is controlled by a neural network and is defined in terms of a sensor mechanism (subsection 2.2.1) and a brain (subsection 2.2.2). The fitness function F of the evolutionary-computation algorithm is given in subsection 2.2.3.

2.2.1 Sensor mechanism

The sensor mechanism of the robot contains a sensor and an internal Gaussian noise source. The sensor detects the nearest food element within its circular field of view with radius r. It only processes the nearest food element within the field of view; the sensor is orientation sensitive. Defining the orientation of the robot by α and the orientation of the nearest food element by β, the sensor activation I (i.e. the input) is given by the normalised one-dimensional Von Mises basis function (Gutowski, 2001).

$$I = \frac{e^{kcos(\alpha-\beta)}}{e^{k}} \tag{1}$$

where k is a positive constant proportional to the width of the basis function. The Von Mises basis function is the spherical analogue of the Gaussian basis function. The normalisation constant e^k ensures that the maximal value of the right-hand side term equals 1 when $\alpha = \beta$.

The internal Gaussian noise source (zero mean, standard deviation sd) is to be considered as a second input: it models the intrinsic noise of neural systems (cf. Zippelius and Engel, 2003). A food element is collected when the distance between the food element and the robot equals $0.1r$.

2.2.2 Brain

The brain (or controller) of the robot is a recurrent neural network with two input neurons (a food-detecting input and a noise source), H hidden neurons, and two output neurons. The input neurons are connected to the hidden neurons and to the output neurons. The hidden neurons have neural inertia by means of self-connections (Nolfi, 2002). This means that the activity of the hidden neurons can be sustained for an extended period of time. The decay in activation is inversely proportional to the weight on the self-connection of the neuron. The initial network configuration is shown in figure 2.

Insert figure 2 about here.

Each connection can be switched on or off during the evolutionary process, while retaining its weight value (cf. Spronck, Sprinkhuizen-Kuyper, and Postma, 2001). Initially, all weights are assigned random values symmetrically distributed around zero on the interval $[-rw,rw]$, with $rw > 0$. The transfer function for the hidden neurons is the sigmoid *tanh* function that maps onto the interval $\langle -1, 1\rangle$. The two output neurons control the robot's relative direction and step size, respectively. The output of the direction neuron is multiplied by π. A modulo operation restricts the direction to the interval $\langle -\pi, \pi\rangle$. The transfer function of the step-size output neuron is a semi-linear function $l = f(u)$ that maps negative values to zero and positive values u to the interval $\langle 0, uL/2\rangle$, with $L = 10\ 000$.

2.2.3 Evolutionary-computation algorithm

The real-valued weights of the neural network controlling the robot are optimised for foraging efficiency using the standard evolutionary-computation techniques of reproduction, crossover, and mutation (Goldberg, 1986, Spronck et al., 2001). The algorithm uses a fitness function

F, which is defined as follows.

$$F = \frac{1}{T}\Big(\sum_{t=1}^{T} c(t) - \lambda \sum_{t=1}^{T} l(t)\Big) \tag{2}$$

where t is an index for individual simulation steps ($t = 1, 2, \ldots, T$) with T denoting the total number of steps, $c(t)$ is a function that returns 1 if a food element is collected at step t and 0 otherwise, $l(t) = f(u, t)$ is the step length of the robot at step t as defined in subsection 2.2.2. The first term between the brackets favours food collection. The second term punishes long steps. The strength of the second term is set by λ.

3 Experiments and results

A large series of experiments was performed to optimise the foraging behaviour of robots. The goal of the experiments was to collect sufficient data of efficient foraging behaviour as required for macroscopic analysis. The experiments are described in subsection 3.1. In subsection 3.2, the resulting foraging behaviour is described.

3.1 Experiments

The foraging behaviour of 1000 robots was optimised. The optimal behaviour of each robot was the result of applying the evolutionary-computation algorithm over 1000 generations with a population size of 1000 robots. All experiments were performed with the following parameter values: H (the number of hidden units) = 2, r (the radius of the circular field of view) = 1, L (the width and heigth of the environment) = 10 000, n (the number of food elements) = 100, rw (the parameter that defines the range in which weights are assigned values) = 0.5, k (the positive constant that is proportional to the width of the basis function) = 20, sd (the standard deviation of the Gaussian noise source) = 0.5, T (the total number of steps) = 10 000, and λ (the parameter that determines the impact of punishment for long steps) = 0.00001.

3.2 Results

The experiments yielded a population of 1000 optimised foragers, with $-0.0045 \leq F \leq 0.0065$. After observing the behaviour of the optimised robots, we distinguished three types of foraging behaviour: (1) systematic behaviour (e.g. foraging in straight lines or circles), (2) random-walk behaviour, and (3) Lévy-flight behaviour. Systematic strategies were employed by 719 of the optimised robots, and considered sub-optimal because these strategies were employed by the optimised agents with the lowest fitness ($F < 0.0$). These 719 robots were left out of all further analyses. The other 281 optimised robots either employed the random-walk strategy shown in figure 3(a) or exhibited the Lévy-flight strategy shown in figure 3(b).

Insert figure 3 about here.

Insert figure 4 about here.

Although most of these 281 fittest optimised robots exhibited the random-walk behaviour, the best ones exhibited Lévy-flight behaviour.

A characteristic feature of the fittest robots (see figure 3(b)) is that their local random-walk behaviour is occasionally interrupted by large jumps. As a result, the area covered by these robots is much larger than the area covered by robots that only exhibit the random-walk behaviour (hence, there is a large difference between the depicted ranges of the x-axis and the y-axis in figures 3(a) and 3(b)). The sudden jumps are known as Lévy flights (Gutowski, 2001, Viswanathan et al., 1999, 2001). Foragers adopting a Lévy-flight strategy outperform the robots adopting a random-walk strategy in terms of foraging efficiency. The advantage of Lévy flights is that the probability of returning to a previously-visited site is smaller than it is for pure random-walk behaviour (Shlesinger and Klafter, 1985).

4 A macroscopic analysis

Inspired by the success of macroscopic analysis in statistical physics, we conducted a macroscopic analysis on the efficient foraging behaviour of the optimised robots. In our context, macroscopic analysis involves the following three main steps: (1) the collection of a very large number of instantiations of the system or behaviour under study, (2) summarising the data in terms of a distribution or descriptors thereof, and (3) the identification of a property that is characteristic for (the behaviour of the) system. Our macroscopic analysis focuses on the quantification of the difference between random walks and Lévy flights in terms of a single parameter μ. The first step of the macroscopic analysis consists of the measurement of step lengths for a very large number of foraging robots. Then, in the second step we summarise the data so obtained by a descriptor (μ) of the shape of the step-length distribution (Viswanathan et al., 1999). Finally, in the third step, we identify the value of μ to be characteristic for the foraging behaviour of the robots. Concentrating on the probability of large step lengths, the tail of the step-length distribution scales according to (cf. Viswanathan et al., 2001):

$$P(l) = \frac{l^{-\mu}}{Z} \tag{3}$$

with $P(l)$ representing the probability of a step of length l, and Z a normalising constant. The parameter μ is proportional to the slope of the tail of the step-length distribution. In particular for a Gaussian function that generates random-walk behaviour, the parameter μ is larger than 3.0. Lévy-flight behaviour is associated with $1.0 < \mu \leq 3.0$. These values of μ yield 'fatter' tails. In our robot, the step-length distribution is generated from the step lengths produced by the output neuron.

To perform our macroscopic analysis we created step-length histograms of the foraging behaviour of optimised robots exhibiting both the random-walk and Lévy-flight behaviours. Figure 4 shows an example of such a step-length histogram.

Insert figure 5 about here.

Subsequently, we analysed the tails of the log-transformed histograms by fitting a linear regression line through the data points. The slope of the line is an estimate of the value of μ that underlies the behaviour of both the random-walk behaviour and the Lévy-flight behaviour of robots. Figure 5 displays the regression line for a robot that exhibits the Lévy-flight behaviour shown in figure 3(b).

Insert figure 6 about here.

The slope of the regression line of the log data is approximately equal to -2 (i.e. $\mu \approx 2$). In terms of equation 3 this corresponds to $P(l) = l^{-2}/Z$.

For each of the 281 fittest foragers a histogram (such as shown in figure 4) was created from several runs of T steps each. Subsequently, the value of μ was determined for each histogram. Figure 6 plots the fitness of the 281 fittest foragers as a function of μ. The values of μ ranged from $\mu \approx 2.0$ to $\mu \approx 4.3$. Figure 6 shows an increase of fitness for agents with a distribution of step lengths that can be characterised by a value of μ that is closer to 2. In other words, the best performing foragers, viz. those employing a Lévy-flight strategy, can be associated with $\mu \approx 2.0$. A range of animals exhibits efficient foraging behaviour that is characterised by Lévy flights with $\mu \approx 2.0$: albatrosses, foraging bumblebees, deer, and amoebas (Viswanathan et al., 1999, 2001).

Insert figure 7 about here.

We provisionally conclude that the macroscopic analysis revealed a property, viz. Lévy-flight behaviour as characterised by $\mu \approx 2.0$, that is universal to optimal foraging behaviour.

5 A microscopic analysis

In addition to the macroscopic analysis, we performed a microscopic analysis to investigate whether the property revealed by the macroscopic analysis can be recognised in the microscopic behaviour of the robots. Our microscopic analysis focused on a successful foraging robot, i.e. a robot that exhibits Lévy-flight behaviour. The analysis aimed at explaining the mechanism that generates the efficient foraging behaviour by inspecting the optimised weights of a successful robot. An interesting example of such a robot is shown in figure 7: the food-detecting sensor of this robot is disconnected.

Insert figure 8 about here.

As can be seen in the figure, the internal noise fully determines the behaviour of the robot. Our analysis of the functioning of this particular robot revealed the following mechanism to underlie the Lévy-flight behaviour. The noise source excites the left hidden neuron randomly. Whenever the excitation exceeds the threshold value of the neuron, it is activated and generates a large step (i.e. a flight). At the same time, its activation excites the right hidden neuron that, in turn, inhibits the left neuron. The neural inertia of the right neuron sustains the inhibition for an extended period. As a result, the large steps occur infrequently, as required for the efficient Lévi-flight strategy. The neural inertia realises a memory that acts as a counting mechanism. Next to this extreme example there were other examples (some with connected food-detecting sensors) of efficient foraging robots.

We provisionally conclude that the microscopic analysis can explain the universal property revealed by macroscopic analysis at a microscopic level. It is important to note that, without macroscopic analysis, microscopic analysis could have revealed Lévy-flight behaviour to be characteristic for optimal foraging, but could not have revealed the exact relation between foraging success and the property μ, and, hence, could not have revealed the universal property of optimal foraging behaviour $\mu \approx 2.0$.

6 Discussion

Below, we discuss the macroscopic and microscopic analyses. In section 6.1, we provide an example of how macroscopic analysis can be applied to complement microscopic analysis in existing research in evolutionary robotics. In section 6.2, we discuss our macroscopic and microscopic analyses of foraging behaviour in relation to the framework of scientific research as described by Cohen (1995).

6.1 Applying macroscopic analysis

To illustrate how macroscopic analysis might generalise findings from existing studies in evolutionary robotics, we adopt the results of a study performed by Nolfi (2002). First, we briefly discuss the study and its findings. Then, we explain how macroscopic analysis can be applied to generalise those findings.

Nolfi (2002) performed a microscopic analysis on agents that had to accomplish a task of self-localisation while driving around in an environment. He investigated a neuro-controller with neurons that change at different time rates. The changing rate, i.e. the neural inertia, of individual neurons was optimised by an evolutionary algorithm. The experiments with the self-localisation task resulted in agents with both fast changing neurons and slow changing neurons. While the fast changing neurons detected regularities occurring at short time scales, the slow changing neurons detected regularities at longer time scales. The first finding of the study is that the self-localisation task requires an agent to be capable of detecting regularities on different time scales. The second finding is that the evolvable changing rates offer this capability to the studied agent.

A macroscopic analysis might be a valuable tool to make both findings easier to generalise. The first finding can be generalised by a macroscopic analysis that relates the self-localisation task to the distribution of neural changing rates of successful agents. It suggests that successful agents should possess both fast and slow changing neurons to be able to extract regularities on short and long time scales. Hence, the above-mentioned distribution might reveal a universal property of self-localisation, i.e. a property characteristic of the distribution of changing rates. To generalise the second finding, a similar macroscopic analysis can be performed on robots performing another task that poses different requirements on the detection of regularities in time. If another task only requires exploitation of regularities on short time scales, then the second finding suggests that evolution will result in robots with neurons that change fast. Hence, we expect the macroscopic analysis of such a task to result in a distribution of changing rates that is different from that characteristic of the self-localisation task.

In summary, the macroscopic analysis that we suggest might be able to reveal a universal property of self-localisation.

6.2 Relation to the framework of scientific research

According to Cohen (1995: 4) the task of science is to provide theories to answer the following three basic research questions. (1) How will a change in the agent's structure affect their behaviour given a task and an environment? (2) How will a change in an agent's task affect its behaviour in a particular environment? (3) How will a change in an agent's environment affect its behaviour on a particular task?

In the context of our study, the agent corresponds to the robot (i.e. the structure consists of the neuro-controller and sensor mechanism; the behaviour is defined by the relative direction and step size provided by the two output neurons), the task is foraging, and the environment consists of randomly distributed food elements. The first basic research question is addressed in our study by evolving the robot's neuro-controller to enhance the foraging behaviour. Both macroscopic and microscopic analyses answer this question, albeit in different ways. The macroscopic analysis answers this question by relating the distribution of step-lengths produced by the robot's neuro-controller to the foraging behaviour observed (see section 4). The microscopic analysis answers this question by relating the structure of the neuro-controller to the observed behaviour of the robot (see section 5). The second basic research question is not addressed in our study because we do not vary the task of foraging. Finally, the third basic research question is addressed by analysing the way in which a robot copes with the change in distribution of food elements in the environment that results from the robot's own foraging. The macroscopic analysis can supply the answer to this question on its own, whereas the microscopic analysis depends on the result from the macroscopic analysis to answer it. The macroscopic analysis answers this question by revealing the universal property of $\mu \approx 2.0$ as the optimal way to cope with the changing distribution of food elements (see section 4). The microscopic analysis cannot answer the third basic research question on its own, because the analysis of a single (or sometimes a few) behaviour(s) cannot detect a universal property that is revealed by the macroscopic analysis. The reason is that the universal property is a characteristic of the average robot-environment interaction, rather than a characteristic of a single (or a few) robot-environment interaction(s). Only after the macroscopic analysis had revealed the answer to the third basic research question (viz. Lévy flights as a characteristic of efficient foraging) the mechanism causing the behaviour at a microscopic level (viz. the neural inertia of the analysed neuro-controller) could be determined. So, macroscopic analysis enabled us to answer the third basic research question. To perform the task of science, viz. provide theories to answer the three basic research questions (Cohen, 1995), macroscopic analysis is preferred over microscopic analysis. Figure 8 illustrates this by depicting the space of basic research questions.

Insert figure 9 about here.

Understanding progresses from descriptions, through prediction, towards causal explanations. Many early microscopic analyses of robots comprise merely descriptions of robot behaviour (see, e.g. Brooks, 1986, Steels, 1994). More recently, microscopic analyses have become more focused on the underlying mechanisms (see, e.g. de Croon et al., in press). Our contribution is to move from specific models towards more general models along the generalisation axis (see figure 8) by employing macroscopic analysis. Evidently, additional macroscopic analyses of robot behaviour are needed to arrive at more general models. As in many other domains where macroscopic analysis has been successfully applied (Stanley et al., 2001), it may lead to a conceptual framework for universality in complex behavioural systems. Future investigations may reveal to what extent other insights in statistical physics can be transferred to robotics.

7 Conclusions

Using macroscopic analysis, we revealed a universal property of foraging behaviour in evolutionary optimised robots, viz. Lévy flights as characterised by $\mu \approx 2$. By doing so, we have shown that macroscopic analysis of a robot model can reveal a universal property of adaptive behaviour.

From the analyses and discussion we may draw the following three conclusions. (1) By combining the two provisional conclusions that followed from our macroscopic and microscopic analyses of efficient foraging behaviour (see sections 4 and 5, respectively) we arrive at the conclusion that macroscopic analysis may predict a universal property that can be explained at the microscopic level by microscopic analysis. (2) From the discussion in subsection 6.1, on an example of macroscopic analysis applied to existing research in evolutionary robotics, we conclude that macroscopic analysis may complement microscopic analysis in the study of adaptive behaviour. (3) From the discussion in subsection 6.2, on how our macroscopic and microscopic analyses of foraging behaviour relate to the framework of scientific research described by Cohen (1995), we conclude that macroscopic analysis may be preferred over microscopic analysis, owing to its power to reveal universal properties.

Given these three conclusions, and the success of macroscopic analysis in statistical physics and other disciplines (Stanley et al., 2001), we expect it to generate novel insights into the universal properties of adaptive behaviour in robot models and natural systems, which is crucial to the understanding of the interaction between a (robot) brain and its environment.

Acknowledgements

We would like to thank the anonymous referees for their constructive comments that helped to improve the paper considerably.

References

Beekman, M., Sumpter, D. J. T. and Ratnieks, F. L. W., 2001, Phase transitions between disordered and ordered foraging in pharaoh's ants. **Proceedings of the National Academy of Sciences**, **98**, 9703-9706.

Beer, R. D., 1995, Computational and dynamical languages for autonomous agents. In *Mind as Motion: Dynamics, Behavior, and Cognition*, edited by R. Port and T. J. van Gelder (Cambridge, MA: MIT Press), pp. 121-148.

Beer, R. D., 2000, Dynamical approaches to cognitive science. **Trends in Cognitive Sciences**, **4**, 91-99.

Beer, R. D., 2003, The dynamics of active categorical perception in an evolved model agent. **Adaptive Behavior**, **11**, 209-243.

Brooks, R. A., 1986, A robust layered control system for a mobile robot. **IEEE Journal of Robotics and Automation**, **2**, 14-23.

Cohen, P. R., 1995, *Empirical methods for artificial intelligence* (Cambridge, MA: MIT Press).

de Croon, G., Nolfi, S. and Postma, E. O., in press, Toward pro-active embodied agents: On the importance of neural mechanisms suitable to process information in time. In *Complex Engineering Systems*, edited by D. Braha, A. Minai and Y. Bar-Yam (Boulder, CO: Perseus Books Group).

Goldberg, D. E., 1986. *Genetic algorithms in search, optimization, and machine learning.* (Reading: Addison-Wesley Publishing Company).

Guillot, A. and Meyer, J., 2001, The animat contribution to cognitive systems research. **Cognitive Systems Research**, **2**, 157-165.

Gutowski, M., 2001, Lévy flights as an underlying mechanism for global optimization algorithms. **Math-ph/0106003**.

Havlin, S., Buldyrev, S. V., Bunde, A., Goldbergera, A. L., Ivanov, P., Peng, C. K. and Stanley, H. E., 1999, Scaling in nature: From DNA through heartbeats to weather. **Physica A**, **273**, 46-69.

Jones, C. and Mataric, M. J., 2003, Adaptive division of labor in large-scale multi-robot systems. IEEE/RSJ international conference on intelligent robots and systems (IROS-03), Las Vegas, NV, 27-31 October 2003.

Lerman, K. and Galstyan, A., 2003, Macroscopic analysis of adaptive task allocation in robots. IEEE/RSJ international conference on intelligent robots and systems (IROS-03), Las Vegas, NV, 27-31 October 2003.

Lerman, K., Galstyan, A., Martinoli, A. and Ijspeert, A. J., 2001, A macroscopic analytical model of collaboration in distributed robotic systems. **Artificial Life**, **7**, 375-393.

Nolfi, S., 1998, Evolutionary robotics: Exploiting the full power of self-organization. **Connection Science**, **10**, 167-184.

Nolfi, S., 2002, Evolving robots able to self-localize in the environment: The importance of viewing cognition as the result of processes occurring at different time scales, **Connection Science**, **14**, 231-244.

Nolfi, S. and Floreano, D., 2000, *Evolutionary Robotics: The Biology, Intelligence and Technology of Self-Organizing Machines* (Cambridge, MA: MIT Press).

Shlesinger, M. F. and Klafter, J., 1985, Lévy walk vs. lévy flights, In *On Growth and Form*, edited by H. E. Stanley and N. Ostrowsky (Dordrecht, The Netherlands: Martinus Nijhoff Publishers).

Slocum, A. C., Downey, D. C. and Beer, R. D., 2000, Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. The Sixth International Conference on the Simulation of Adaptive Behavior, 11-15 September 2000 (Cambridge MA: MIT Press), pp. 430-439.

Spronck, P. H. M., Sprinkhuizen-Kuyper, I. G. and Postma, E. O., 2001, Evolutionary learning of a neural robot controller. Proceedings of the International Conference on Computational Intelligence for Modelling, Control and Automation, 9-11 July 2001 (Vienna, Austria: IOS Press), pp. 510-518.

Stanley, H. E., Amaral, L. A. N., Gopikrishnan, P., Ivanov, P. C., Keitt, T. H. and Plerou, V., 2001, Scale invariance and universality: Organizing principles in complex systems. **Physica A**, **281**, 60-68.

Steels, L., 1994, Emergent functionality in robotic agents through on-line evolution. The 4th Artificial Life conference (alife iv), 6-8 July 1994 (Cambridge, MA: MIT Press), pp. 8-16.

van Dartel, M. F., Sprinkhuizen-Kuyper, I. G., Postma, E. O. and van den Herik, H. J., in press, Reactive agents and perceptual ambiguity. **Adaptive Behavior**.

Viswanathan, G. M., Afanasyev, V., Buldyrev, S. V., Havlin, S., Luz, M. G. E. da Raposo, E. P. and Stanley, H. E., 2001, Lévy flights search pattern of biological organisms. **Physica A**, **295**, 85-88.

Viswanathan, G. M., Buldyrev, S. V., Havlin, S., Luz, M. G. E. da, Raposo, E. P. and Stanley, H. E., 1999, Optimizing the success of random searches. **Nature**, **401**, 911-914.

Zippelius, A. and Engel, A., 2003, Statistical mechanics of neural networks. In *The handbook of brain theory and neural networks* edited by Arbib, M. (Cambridge, MA: MIT Press), pp. 930-934.

8 Figure Captions

Figure 1. Illustration of a robot (circle) in the experimental environment with randomly distributed food elements (dots). The values on the x-axis and y-axis are spatial coordinates ($0 \leq x, y < L = 10\ 000$).

Figure 2. The initial network configuration of robots entering the evolutionary optimisation process.

Figure 3. Two types of foraging behaviour.

Figure 4. Histogram of step lengths.

Figure 5. Log-log plot of the tail (l >> 0) of the histogram. The slope of the regression line is ≈ -2.

Figure 6. The fitness F (see section 2.2.3) as a function of the step-size distribution parameter μ for the 281 fittest foragers obtained through evolutionary optimisation of foraging behaviour.

Figure 7. An example of the network configuration of a successful foraging robot.

Figure 8. The two-dimensional space of versions of the basic research questions (Redrawn from Cohen, 1995).

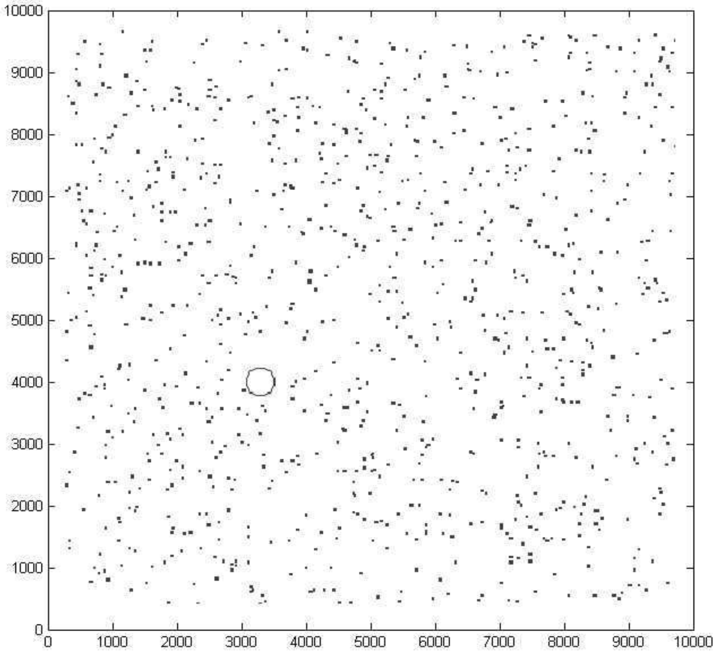
10000
9000
8000
7000
6000
5000
4000
3000
2000
1000
0
0
1000
2000
3000
4000
5000
6000
7000
8000
9000
10000

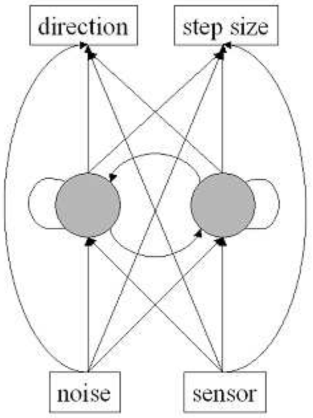
direction
step size
noise
sensor

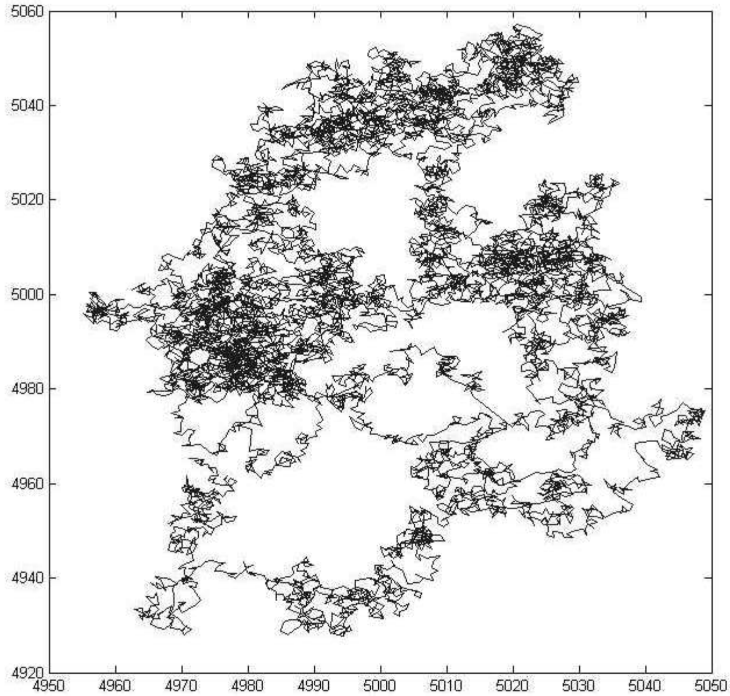

(a) Illustration of random-walk foraging. The path consist of 10 000 steps. The values on the x-axis and y-axis are spatial coordinates with $0 \leq x, y < L = 10\ 000$, of which only the ranges $4950 \leq x < 5050$ and $4920 \leq y < 5060$ are depicted

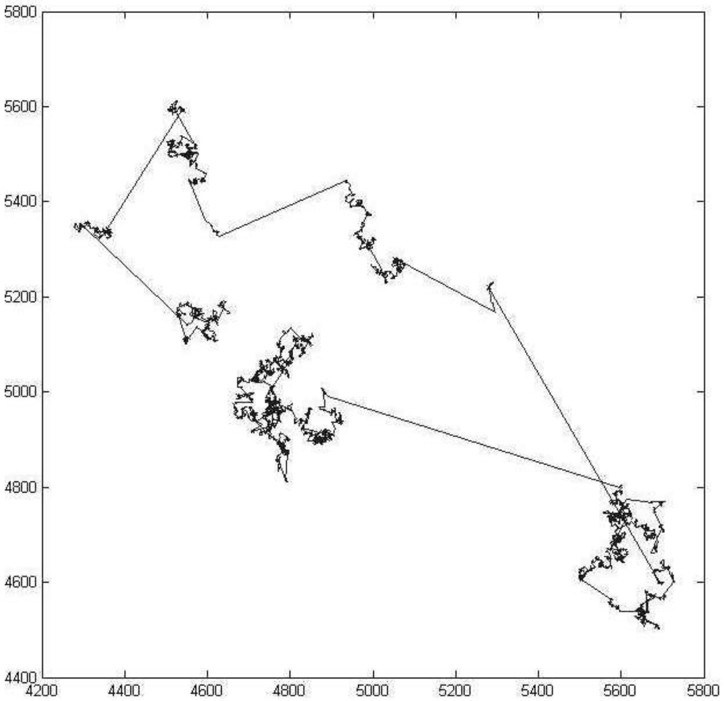

(b) Illustration of Lévy-flight foraging. The path consists of 10 000 steps. The values on the x-axis and y-axis are spatial coordinates with $0 \leq x, y < L = 10\ 000$, of which only the ranges $4200 \leq x < 5800$ and $4400 \leq y < 5800$ are depicted

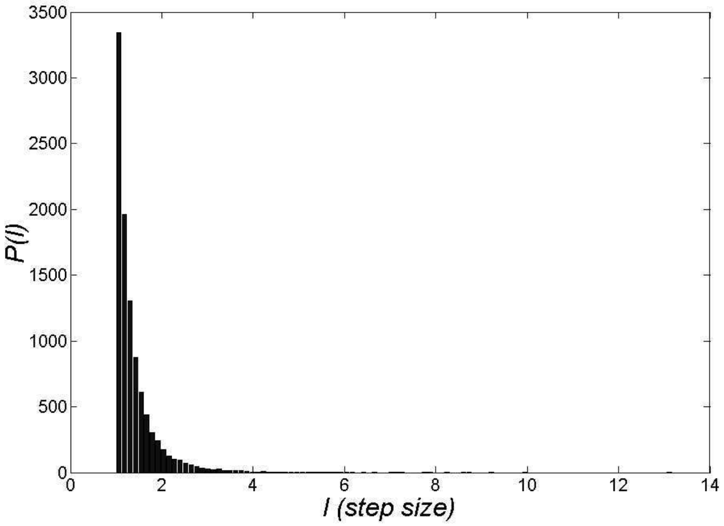
3500
3000
2500
2000
1500
1000
500
0
P(l)
0
2
4
6
8
10
12
14
l (step size)

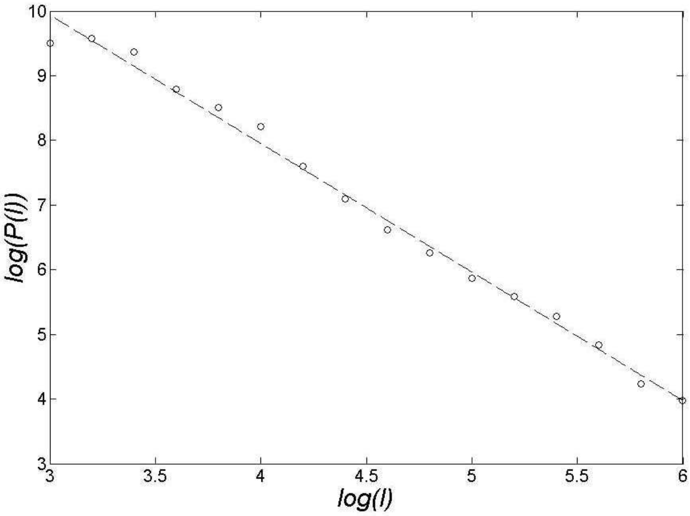
10
9
8
7
6
5
4
3
log(P(l))
3
3.5
4
4.5
5
5.5
6
log(l)

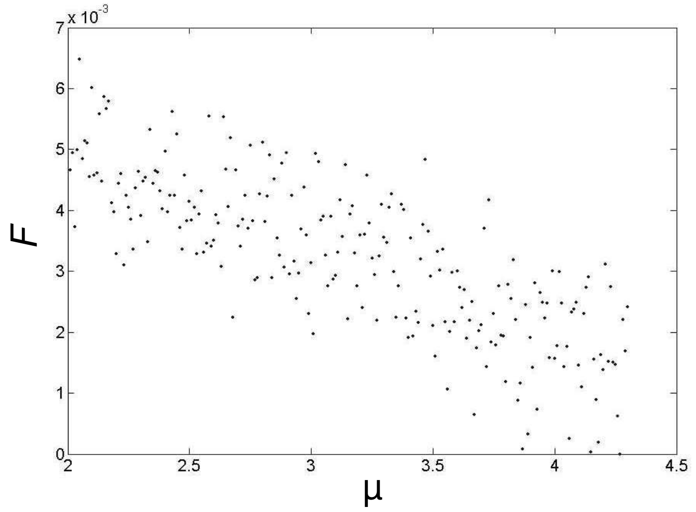

x 10^-3
7
6
5
4
3
2
1
0
F
2
2.5
3
3.5
4
4.5
μ

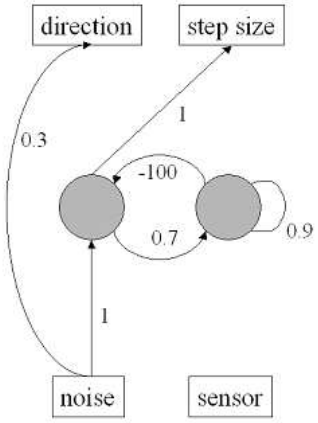
direction
step size
0.3
1
-100
0.7
0.9
1
noise
sensor

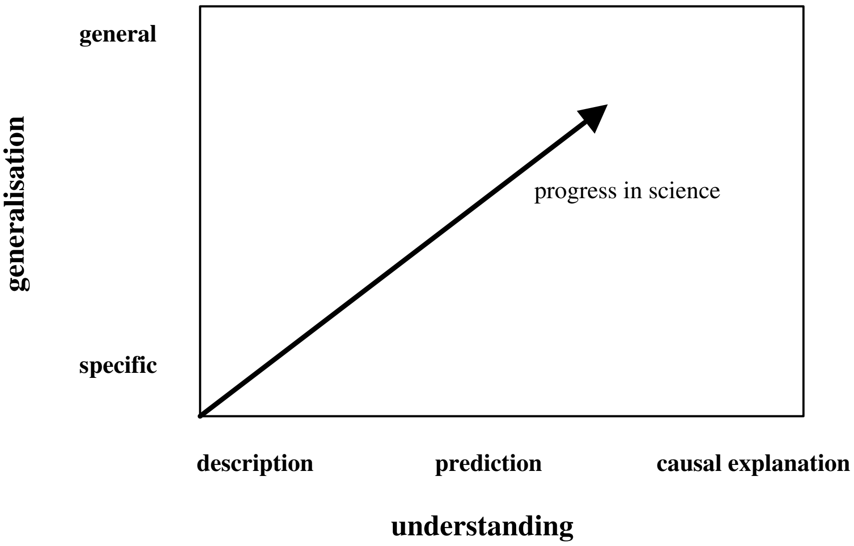
general
generalisation
specific
progress in science
description
prediction
causal explanation
understanding