

On the Role of Quantitative Descriptions of Behaviour in Mobile Robotics Research

Ulrich Nehmzow

Dept. of Computer Science, University of Essex, Colchester CO4 3SQ, UK

Abstract. This paper – a summary of a keynote address given at the Robocup 2003 symposium – argues i) that mobile robotics research would benefit from a *theoretical* understanding of robot-environment interaction, ii) that independent replication and verification of experimental results should become common practice within robotics research, and iii) that *quantitative* measures of robot behaviour are needed to achieve this.

The paper gives one example of such quantitative measures of behaviour: the reconstruction of the phase space describing a robot's behaviour, and its subsequent analysis using chaos theory.

1 Mobile Robotics Research

1.1 Robot Engineering versus Robot Science

Theory Supports Design. Arguably, there are (at least) two independent objectives of robotics research: on the one hand, to create artefacts that are capable of carrying out useful tasks in the real world – for example industrial, service, transportation or medical robots, to name but a few, and on the other hand to obtain a theoretical understanding of the design issues involved in making those artefacts – for example sensor and actuator modelling, system identification (modelling of entire systems), or sensor, actuator and behaviour analysis. The former can be referred to as ‘robot engineering’, the latter as ‘robot science’. It is robot science that this paper is concerned with.

While robot engineering ultimately produces the ‘useful’ artefacts, there is a lot that robot science can contribute to this process. Without theoretical understanding, any design process is largely dependent upon trial-and-error experimentation and iterative refinement. In order to design in a principled way, a hypothesis of some kind (a justified expectation) is needed to guide the design process. The hypothesis guides the investigation: results obtained are fed back into the process and brought into alignment with the theory, to lead to the next stage of the experimentation and design. The better the theory underlying the design process, the more effective and goal-oriented the design process will be.

Every process of designing technical artefacts is based on some kind of assumptions (a ‘theory’), even if very little is known at all about the object being designed. This is true for current mobile robotics research, too. When asked to design a wall-following robot, the designer will not start with an *arbitrary*

program, but with a ‘reasonable guess’, sensibly speculating on which sensors might be useful to achieve the desired behaviour, which general kind of control program will perform acceptably, *etc.* But, given our current understanding of robotics, he is unable to design the entire behaviour off-line!

Instead, mobile robotics researchers to-date are crucially dependent on trial-and-error procedures. A ‘reasonable prototype’ has to be tested in the target environment, and refined based on observations and underlying theory (‘hunch’ is often the more appropriate term for such theories). Here is a practical example: to design the Roomba commercial robot floor cleaner (relying on very simple sensing, and not involving any sophisticated navigation), thirty prototypes had to be built over a period of twelve years [EXN 03]!

The first argument we would make in favour of a better theoretical understanding of the principles underlying a mobile robot’s operation in its environment, therefore, is that robot engineering (the process of designing a technical artefact that will perform useful tasks in the real world) will benefit from theory through the resulting more effective, rigorous and goal-oriented development methods.

Science Requires Replication and Verification. Current mobile robotics research practice not only differs from that of established disciplines in its lack of theories supporting design, but also in a second aspect: independent replication and verification of experimental results is uncommon. While in sciences such as biology or physics, for instance, reported results are only taken seriously once they have been verified independently a number of times, in robotics this is not the case. Instead, papers often describe experimental results obtained in specific environment, under specific experimental conditions. These experiments therefore are ‘existence proofs’ – the demonstration that a particular result can be achieved – but they do not state in general terms under which conditions a particular result can be obtained, nor which principles underlie the result. Existence proofs are useful, they demonstrate that something can be achieved, which is an important aspect of science, but they do not lead towards general principles and theories.

The second argument we make, therefore, is that mobile robotics research is now at a stage where we should move on from existence proofs to a research culture that habitually includes independent replication and verification of experiments.

The Role of Quantitative Descriptions. Theories, experimental replication and experimental verification all depend crucially on *quantitative* descriptions: quantitative descriptions are an essential element of the language of science.

The third argument we make, therefore, is that an essential first step towards a sounder theoretical understanding of robot-environment interaction is to develop and apply quantitative descriptions of robot-environment interaction. The experiments reported in this paper are one example of how to achieve this.

1.2 Theory

Introduction. When referring to ‘theory’, we mean a coherent body of hypothetical, conceptual and pragmatic generalisations and principles that form the general frame of reference within which mobile robotics research is conducted.

There are two key elements that make a theory of robot-environment interaction useful, and therefore desirable for research:

1. A theory will allow the formulation of hypotheses for testing. This is an essential component in the conduct of ‘normal science’ [Kuhn 70].
2. A theory will make predictions (for instance regarding the outcome of experiments), and thus serve as a safeguard against unfounded or weakly supported assumptions.

A theory retains, in abstraction and generalisation, the essence of what it is that the triple of robot-task-environment does (see also figure 2). This generalisation is essential: it highlights the important aspects of robot-environment interaction, while suppressing unimportant ones. Finally, the validity of a theory (or otherwise) can then be established by evaluating the predictions made applying the theory.

Benefits of Theory. Two significant advantages of a theory have been given above: generating hypotheses and making testable predictions. But there are practical advantages, too, particularly for a discipline that involves the design of technical artefacts: Theory supports off-line design, i.e. the design of technical artefacts through the use of computer models, simulations and theory-based calculations.

Example: Aircraft Design. Because of a theoretical understanding of aerodynamics, it is now possible to design aircraft (such as the Airbus A380, shown in figure 1) almost completely ‘off-line’. This incurs considerable advantages.

In a report on ‘virtual manufacturing’ to the UK Foresight Programme, J. Coyle states that virtual manufacturing and virtual prototyping – validating product design and production processes in a synthetic environment – enabled aerospace manufacturing companies to reduce design to first-unit build time by 33%, using 25% less manpower. In addition to that, the amount of rework is reduced by 90% (Boeing claim), sustained engineering work is reduced by 50%, and production time cycles and planning costs are also reduced. In aircraft design, there is the additional advantage that simulations are reusable for operator work instructions and for maintenance tasks.

These considerable benefits are, obviously, based on sound theoretical knowledge of the governing laws of physics, mechanics, aerodynamics, material science *etc.* underlying the operation of the designed artefact.

In mobile robotics, such theoretical knowledge to construct sufficiently accurate computer models and computer simulations is not yet available, but it is obvious that robot design would benefit significantly if it were.



Fig. 1. Aircraft design is largely based on theory, and computer modelling and considerations based on theoretical understanding have reduced development time and cost considerably.

2 Quantitative Characterisations of Robot-Environment Interaction

A key element of the objectives outlined above are *quantitative* descriptions of the robot's behaviour. They will support replication and independent verification of experimental results, principled experimental design and analysis of the influence of individual experimental factors. In this paper we present a way of using methods from deterministic chaos and dynamical systems theory to achieve this.

2.1 Robot-Environment Interaction Constitutes a Dynamical System

The behaviour of a mobile robot is governed by three main factors: (i) the robot itself: its hardware, physical makeup, inertia, sensor and motor characteristics *etc.*; (ii) the environment the robot is operating in: the colour, texture and structure of walls, floors and ceilings, the temperature and humidity, speed of sound in the environment, noise *etc.*; and (iii) the control program (the 'task') the robot is executing (figure 2).

Therefore, we argue, a mobile robot, interacting with its environment, could be viewed as an *analog computer*.

Similar to an optical lens, which takes light rays as its input and 'computes' the Fourier transformation of the input image as its output (thus acting as an analog computer), or a cylindrical lens taking the visual image of the environment as its input and 'computing' the positions of vertical edges in the image as its output (again acting as an analog computer), the mobile robot, executing some control program in its environment 'computes' behaviour as its output from the three input components shown in figure 2, i.e. robot-specific components, environment-specific components and the task (see figure 3).

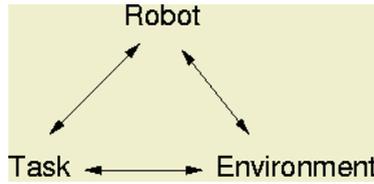


Fig. 2. The fundamental triangle of robot-environment interaction: A robot’s behaviour always has to be seen in the context of robot, task and environment.

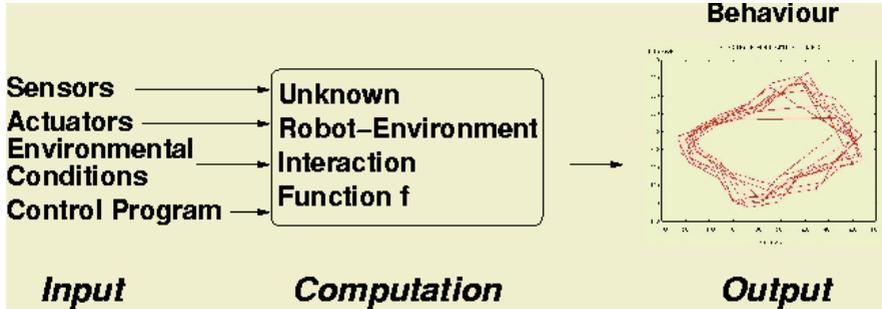


Fig. 3. A mobile robot interacting with its environment can be described as an analog computer, taking environmental, morphological and task-related data as input, and ‘computing’ behaviour as output (see also figure 2).

We furthermore argue that one of the most fundamental manifestations of the analog computation carried out by a mobile robot interacting with its environment, is the trajectory taken by the robot. Although the trajectory does not encompass *every* aspect of a mobile robot’s behaviour, it is the dominant result of the robot’s perception, control architecture and actuation. To analyse the trajectory, therefore, means analysing the dominant aspect of the ‘output’ of the robot’s ‘computation’. (Note: the methods described in this paper have been applied to the analysis of mobile robot behaviour. However, it is conceivable that other kinds of robot behaviour can be analysed in the same way, as our analysis is a time series analysis, which can be used to analyse *any* time series describing a robot’s behaviour, not just trajectories.)

The trajectory of a mobile robot essentially constitutes two time series – one for the robot’s *x*-coordinate, and one for the *y*-coordinate. In our experiments we analyse these time series for the presence or absence of deterministic chaos, and use the quantitative descriptors of deterministic chaos as quantitative descriptions of the robot’s interaction with its environment.

2.2 Chaos Theory to Characterise Robot-Environment Interaction Quantitatively

Introduction to the Approach. Our aim, then, is to obtain a quantitative description of a mobile robot’s interaction with its environment. We achieve

this by applying methods from dynamical systems theory: We first reconstruct the attractor underlying the robot’s behaviour, then analyse the attractor using chaos theory, and describe it quantitatively by computing the Lyapunov exponent (section 2.3). In this way we are able to establish, for instance, that robot environment interaction does exhibit deterministic chaos, as well as the influence of individual experimental parameters (such as objects present in the environment) upon the overall robot behaviour.

Analysis of a Dynamical System in Phase Space. The behaviour of any dynamical physical system is fully described by giving its position $p(t)$ and momentum $m(t)$ at time t , for every degree of freedom this system has. This (p, m) space is the *phase space* of the system – if it is known, the motion of the system is known and analysable¹.

For actual physical systems, the exact nature of the system’s phase space cannot be known, due to noise *etc.* It is, however, possible to reconstruct an approximation of the system’s phase space very simply from an observed time series of the system’s motions, through the so-called time-lag-embedding Y_n given in equation 1 [Peitgen et al 92, Kantz & Schreiber 97, Abarbanel 96].

$$Y_n = (x(t_n - (p - 1)h), x(t_n - (p - 2)h), \dots, x(t_n - h), x(t_n)), \quad (1)$$

with $x(t)$ being a sequential set of measurements (the time series), p being the embedding dimension and h being the embedding lag (for a full discussion see [Peitgen et al 92, Nehmzow & Walker ny]).

In other words, it is not necessary to have full knowledge of a physical system’s phase space, a (sufficiently long) observation of the system’s behaviour (for instance by logging a trajectory, using a camera, or by logging other data describing the agent’s behaviour) is sufficient to reconstruct its phase space. Once reconstructed, the phase space can be analysed quantitatively (see section 2.3).

In summary: We are interested in analysing some physical system quantitatively. Logging data emanating from the system (such as a robot trajectory) for a reasonably long period of time, we obtain a time series describing the system’s behaviour. We know that it is possible to reconstruct the system’s phase space from that time series [Takens 81]: instead of analysing the system’s motion in physical space, we analyse the system’s motion on the ‘attractor’ in phase space. Finally, attractors (and therefore the system that has generated the attractor) can be described quantitatively, for instance through the Lyapunov exponent or the dimension of the attractor [Peitgen et al 92].

Example from Mobile Robotics. Figure 4 shows the motion of a Pioneer II mobile robot in our laboratory, executing an obstacle-avoidance control program

¹ A very simple illustration of phase space is the ideal pendulum, which has one degree of freedom – the arc along which it swings. The phase space of the ideal pendulum is therefore two-dimensional (speed and position along the arc). The actual trajectory of the pendulum through phase space – the ‘attractor’ – is a circle.

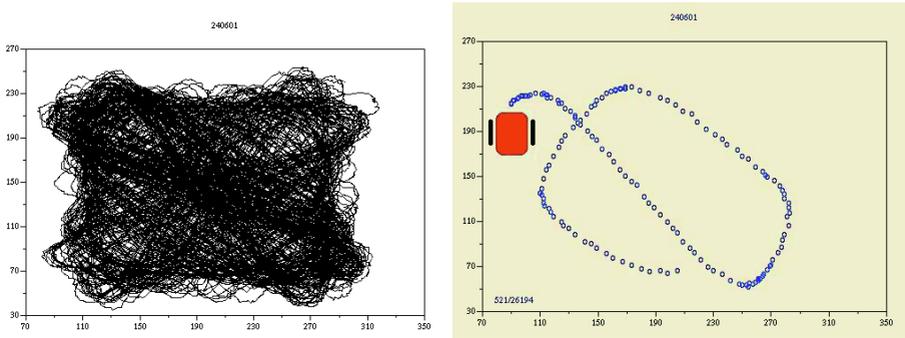


Fig. 4. Quasi ‘Billiard Ball’ Behaviour in Square Arena – Entire trajectory (left) and 150 data points (right).

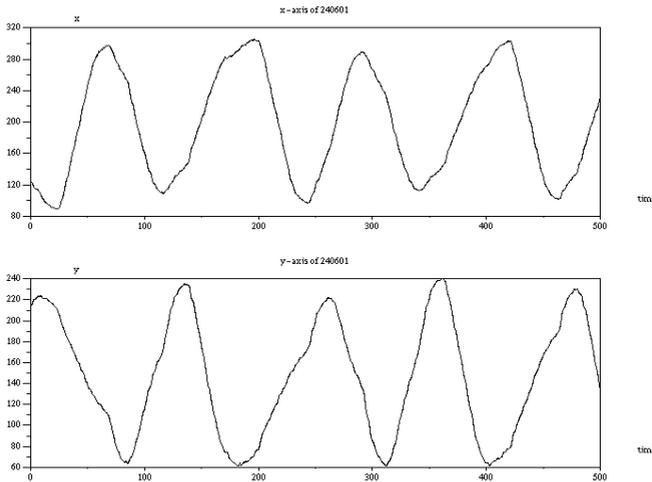


Fig. 5. Part of the x (top) and y (bottom) coordinate of quasi billiard ball behaviour. The attractor describing the robot’s behaviour (shown in figure 6) was reconstructed from this data, using time-lag embedding [Peitgen et al 92].

that results in a quasi billiard ball behaviour. The trajectory shown in figure 4 was obtained by logging the robot’s motion with an overhead camera for about 2 hours of continuous operation.

Using the time series $x(t)$ or $y(t)$, shown in figure 5, the robot’s motion in phase space can be reconstructed through time lag embedding (equation 1). This phase space reconstruction is shown in figure 6.

As the phase space of a physical system describes that system’s behaviour fully, we now have a representation of a robot’s interaction with its environment that can be analysed quantitatively – one of the key requirements of a theory of robot-environment interaction, as we argued earlier. The following section will discuss how the Lyapunov exponent can be used to achieve this objective.

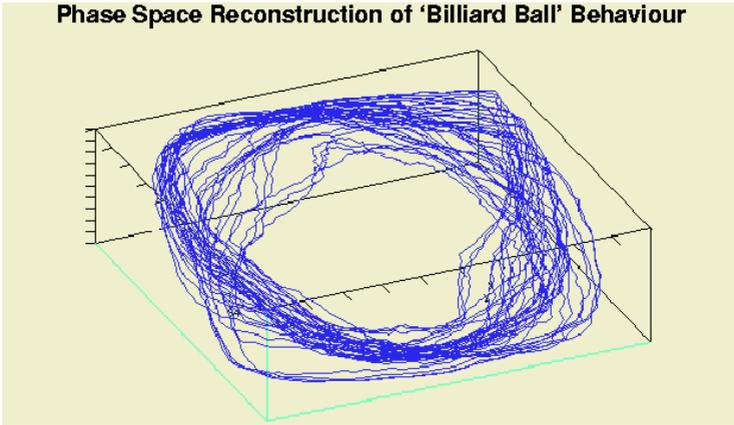


Fig. 6. The phase space reconstruction of the quasi billiard ball behaviour shown in figure 4. The attractor has a fractal dimension of approx. 1.9, it is 'strange'.

2.3 Quantitative Analysis of Phase Space

There are a number of quantitative descriptions of phase space, for instance the dimension of the attractor (correlation dimension), but the description most easily obtained, and which we have used most in our experiments is the Lyapunov exponent.

The Lyapunov Exponent. One of the most distinctive characteristics of a chaotic system is its sensitivity to a variation in the system's variables: two trajectories in phase space that started close to each other will diverge from one another as time progresses, the more chaotic the system, the greater the divergence.

Consider some state S_o of a deterministic dynamical system and its corresponding location in phase space. As time progresses the state of the system follows a deterministic trajectory in phase space. Let another state S_1 of the system lie arbitrarily close to S_o , and follow a different trajectory, again fully deterministic. If d_o is the initial separation of these two states in phase space at time $t = 0$, then their separation d_t after t seconds can be expressed as $d_t = d_o e^{\lambda t}$.

Or, stated differently, consider the average logarithmic growth of an initial error E_0 (the distance $|x_0 - (x_0 + \epsilon)|$, where ϵ is some arbitrarily small value and x_0 a point in phase space) [Peitgen et al 92, p. 709]. If E_k is the error at time step k , and E_{k-1} the error at the previous time step, then the average logarithmic error growth can be expressed by equation 2.

$$\lambda = \lim_{n \rightarrow \infty} \lim_{E_0 \rightarrow 0} \frac{1}{n} \sum_{k=1}^n \log \left| \frac{E_k}{E_{k-1}} \right|. \quad (2)$$

λ (which is measured in Hz or s^{-1} , or sometimes in bits/s) is known as the Lyapunov exponent.

For a m -dimensional phase space, there are m λ values, one for each dimension. If any one or more of those components are positive, then the trajectories of nearby states diverge exponentially from each other in phase space and the system is deemed chaotic. Since any system's variables of state are subject to uncertainty, a knowledge of what state the system is in can quickly become unknown if chaos is present. The larger the positive Lyapunov exponent, the quicker knowledge about the system is lost. One only knows that the state of the system lies somewhere on one of the trajectories traced out in phase space, i.e., somewhere on the strange attractor.

The Lyapunov exponent is one of the most useful quantitative measures of chaos, since it will reflect directly whether the system is indeed chaotic, and will quantify the degree of that chaos. Also, knowledge of the Lyapunov exponents becomes imperative for any analysis on prediction of future states.

One method to determine the Lyapunov of an attractor describing the behaviour of a physical system is to estimate it from an observed time series of the system's motion [Peitgen et al 92]. The methods we have used in our research are the one proposed by Wolf in 1985 [Wolf et al 95], and the one proposed by Abarbanel [Applied Nonlinear Sciences 03]. Software packages implementing either method are readily available.

3 Experiments

3.1 Experimental Setup

In our experiments, we observed autonomous mobile robots (figure 7), performing simple sensor-motor tasks such as obstacle avoidance, in a laboratory environment (see also figure 8).

The robot's trajectory was logged every 250 ms, using an overhead camera. Figures 4 and 5 show examples of the kind of trajectories observed, and the kind of data subsequently analysed.

Results. We carried out a number of different experiments, in which our robots performed different behaviours, in a range of different environments. As this paper is concerned with method, we only give an overview of results here, full results are given in [Nehmzow & Walker 03, Nehmzow & Walker 03b] [Nehmzow & Walker ny].

First, all our experiments showed that the interaction of our mobile robots with their environment exhibited deterministic chaos. Lyapunov exponents were small, but always positive (indicating the presence of chaos), and typically between 0.1 and 0.3 bits/s. All attractors were 'strange', i.e. had a fractal dimension of typically between 1.5 and 2.5.

Second, we were able to investigate the influence of individual experimental parameters quantitatively. For instance, changing the robot's control program



Fig. 7. The Magellan Pro and Pioneer II mobile robots used in the experiments discussed here.

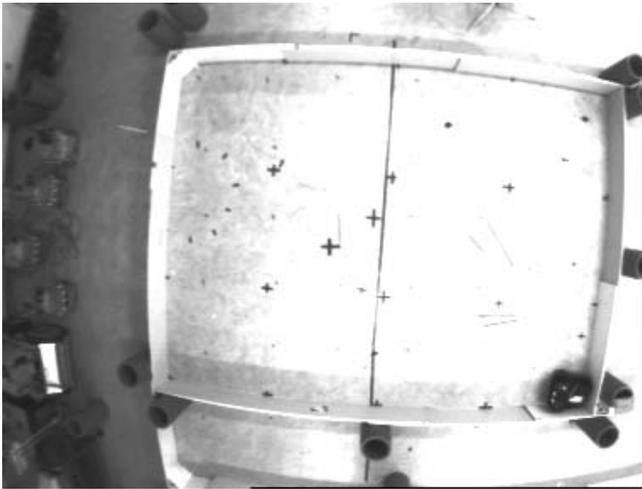


Fig. 8. Bird's eye view of the experimental arena used in our experiments. The robot is visible in the bottom right hand corner of the arena.

from 'wall following' to 'billiard ball behaviour' increased the Lyapunov exponent by a factor of four, whereas changing the environment (by adding a central obstruction to it) whilst leaving the behaviour (billiard ball) unchanged, resulted in no measurable change of the Lyapunov exponent [Nehmzow & Walker ny] [Nehmzow & Walker 03]. In other words: in the experiments we conducted, the control code influenced robot-environment interaction much more noticeably than the nature of the environment!

In a different set of experiments we investigated differences between the Pioneer II and the Magellan Pro robots. Using similar control programs for both robots (the Pioneer and the Magellan use different operating systems, and code can therefore not be *identical*), we found in preliminary experiments that there is a noticeable difference between these two robots: the Magellan attained Lyapunov exponents that were smaller by about a factor of 2. This observation is subject to ongoing research.

4 Summary and Conclusion

4.1 Summary

In this keynote paper we argue that i) the design of technical artefacts (engineering) benefits from theory, that ii) mobile robotics research has reached a stage of maturity that allows the field to move on from existence proofs to a research culture that habitually involves independent replication and verification of experimental results, and that iii) quantitative descriptions of behaviour are the foundation of theory, replication and verification.

An engineering design process that is not guided by a theoretical understanding of the subject matter has to resort to trial-and-error methods and iterative refinement. This is very costly, and does not even guarantee optimal results. On the other hand, design supported by theory can use the hypotheses and predictions afforded by the theory to reduce the design space, resulting in more efficient and faster design cycles. The aircraft industry, for instance, reports substantial gains in efficiency through ‘virtual manufacturing’, based on a theoretical understanding of aircraft design [Coyle 03].

Mobile robotics research, so far, largely relies on trial-and-error methods and uses iterative refinement techniques to develop task-achieving robot controllers. Existing computer models of robot-environment interaction are such simplified representations of sensors, actuators and environments that their predictive power is of little value for real world applications. Instead of mainly designing robot controllers off-line, control programs have to be developed through iterative refinement processes, which require large amounts of time and are costly. Furthermore, because robot control programs are usually developed in specific target environments, experimental results are typically existence proofs, rather than generally applicable findings. This limits their usefulness.

Finally, mobile robotics research currently does not benefit from a research practice in which experimental results are replicated and independently verified. This, of course, increases the risk that results stemming from accidental singularities go undetected.

We argue that one of the reasons for the current practice of mobile robotics research is that we simply haven’t got the means to communicate results in such a way that theoretical descriptions, independent replication and verification are possible – we lack the ‘language’, i.e. quantitative descriptions of robot behaviour. Statistical descriptions are useful in this regard to some extent, but they still only represent a statistical description of *overall* robot behaviour, rather than a quantitative description of a particular single run.

In this paper, we present one method to obtain a quantitative description of robot behaviour, using dynamical systems theory. We first reconstruct the attractor underlying the robot's behaviour, then analyse the attractor using chaos theory, and describe it by computing the Lyapunov exponent (section 2.3). In this way we were able to establish, for instance, that robot environment interaction does exhibit deterministic chaos, as well as the influence of individual experimental parameters (such as objects present in the environment) upon the overall robot behaviour.

4.2 Conclusion

The experiments presented in this paper and elsewhere [Nehmzow & Walker 03] [Nehmzow & Walker ny, Nehmzow & Walker 03b] demonstrate that dynamical systems theory can be used to obtain quantitative descriptions of robot-environment interaction. Computing the Lyapunov exponent of the phase space underlying the robot's behaviour provides a means of measuring the influence of individual experimental parameters upon the overall robot behaviour. By changing just one of the three components shown in figure 2, for instance, the Lyapunov exponent can be used to describe the robot itself, the environment it is acting in, or the task it is performing. Likewise, it can be used to describe the influence of parameters such as illumination, colour, surface structure *etc.* by modifying the parameter in question in a principled way, and measuring the resulting Lyapunov exponent.

Acknowledgements

Much of the work reported in this paper was conducted in collaboration with my colleague Keith Walker from Point Loma Nazarene University, San Diego, I gratefully acknowledge his contribution.

The research presented in this paper was supported by a grant from the Research Promotion Fund of the University of Essex.

I thank the RoboCup organising committee, which invited me to present a keynote address at the 2003 RoboCup symposium in Padua, subsequently leading to this paper.

Thank You also to Phillip McKerrow and Claudia Nehmzow for detailed and helpful comments on earlier versions of this article.

References

- [Abarbanel 96] H. Abarbanel, *Analysis of Observed Chaotic Data*, Springer-Verlag, New York 1996.
- [Applied Nonlinear Sciences 03] Applied Nonlinear Sciences, *Tools for Dynamics*, URL=<http://www.zweb.com/apnonlin>, last accessed May 2003.
- [EXN 03] EXN.CA discovery channel, 11.3.2003, <http://www.exn.ca/Stories/2003/03/11/57.asp>, last accessed October 2003.

- [Coyle 03] J. Coyle, Virtual Manufacturing, UK Foresight Programme Materials and Structures National Advisory Committee, *Aerospace 2020*, AR360 Vol. III, http://www.iom3.org/foresight/nac/html_docs/AGARD%203.htm, Last accessed October 2003.
- [Kaplan & Glass 95] D. Kaplan and L. Glass, *Understanding Nonlinear Dynamics*, Springer Verlag London, 1995
- [Kantz & Schreiber 97] H. Kantz and T. Schreiber, *Nonlinear Time Series Analysis*, Cambridge University Press, 1997.
- [Kuhn 70] Thomas Kuhn, *The Structure of Scientific Revolutions*, University of Chicago Press, 1970.
- [Nehmzow & Walker ny] Ulrich Nehmzow and Keith Walker, Quantitative Description of Robot-Environment Interaction using Chaos Theory, submitted to *J Robotics and Autonomous Systems*.
- [Nehmzow & Walker 03] Ulrich Nehmzow and Keith Walker, The Behaviour of a Mobile Robot Is Chaotic, *AISB Journal*, Vol 1 No 4, 2003.
- [Nehmzow & Walker 03b] Ulrich Nehmzow and Keith Walker, Quantitative Description of Robot-Environment Interaction Using Chaos Theory, *Proc. European Conference on Mobile Robotics (ECMR)*, Warsaw 2003.
- [Peitgen et al 92] H.-O. Peitgen, H. Jürgens and D. Saupe, *Chaos and Fractals*, Springer Verlag 1992.
- [Takens 81] F. Takens, *Detecting Strange Attractors in Turbulence*, Lecture Notes in Mathematics, Vol. 898, Springer Verlag, Berlin 1981.
- [Wolf et al 95] A Wolf, J. Swift, H. Swinney and J. Vastano, Determining Lyapunov exponents from a time series, *Physica 16D*, 1995. Code available at <http://www.cooper.edu/~wolf/chaos/chaos.htm>, last accessed November 2003.