

Towards cognitive robotics : robotics, biology and developmental psychology

LEE, Mark, NEHMZOW, Ulrich and RODRIGUES, Marcos
<<http://orcid.org/0000-0002-6083-1303>>

Available from Sheffield Hallam University Research Archive (SHURA) at:
<http://shura.shu.ac.uk/5282/>

This document is the author deposited version. You are advised to consult the publisher's version if you wish to cite from it.

Published version

LEE, Mark, NEHMZOW, Ulrich and RODRIGUES, Marcos (2012). Towards cognitive robotics : robotics, biology and developmental psychology. In: MCFARLAND, David, STENNING, Keith, MCGONIGLE, Maggie and HENDRY, Derek, (eds.) The complex mind : an interdisciplinary approach. Basingstoke, Palgrave Macmillan, 103-126.

Copyright and re-use policy

See <http://shura.shu.ac.uk/information.html>

Towards Cognitive Robotics: Robotics, Biology and Developmental Psychology

Mark Lee, Ulrich Nehmzow and Marcos Rodrigues

1 Introduction — Early Robotics Research

The question of how “intelligent” controllers can be designed for machines has always attracted much interest and research activity. The idea of reproducing some facets of human cognition in a designed artifact has fascinated scientists, philosophers and charlatans throughout history. But despite the enormous efforts that have been directed at this issue in the twentieth century, only very recently has any significant progress been made. The robots of the last century typically were brittle, they failed in even simple tasks as soon as these tasks deviated even slightly from the original specifications; they were slow; and they needed constant attention from software engineers.

Up to the late 1980s and early 1990s robotics research was dominated by work originating from the control theory and cybernetics communities, which meant that the fundamental assumptions of how intelligent behaviour could be achieved were very similar if not identical. The *Sense-Think-Act* cycle [Nehmzow, 2003] was one such assumption: sensor signals (“sense”) would be perceived by the robot, processed through various processing stages (“think”) and result in motor action (“act”). This cycle would then be repeated. Another common assumption was the Physical Symbol System Hypothesis [Newell and Simon, 1976], which claimed that general intelligent action could be achieved by manipulating symbols (which represented states of the real world) through a set of “the right kind” of rules. Assumptions like these influenced much of early robotics research towards a “programming approach” to intelligent behaviour, involving rules, knowledge-bases, agents and simulated environments.

This changed drastically when Rodney Brooks' introduced the concept of “Behaviour-based Robotics”, and many laboratories started to pursue research in the “New AI” or “Embodied Intelligence” as it became known. This approach considered work with physically embedded machines as essential, used little or no symbols (using neural or “subsymbolic representations”), and saw the behaviour of a robot as emergent from the robot's interaction with the environment, the robot's morphology, and the many unknown factors that influence robot behaviour. In the last decade this approach has blossomed into the “Embodiment” movement which argues that truly autonomous intelligent agents must be situated, embedded, and embodied, and, currently, the only exemplars are to be found in the natural world. This has spurred much biologically-inspired robotics research that has taken ideas and models from brain science (neurology, anatomy, physiology), psychology (behaviour, perception and psychophysics), cognitive science, ethology, and even evolutionary theory.

This chapter has its origins in an early collaborative robotics research project conducted jointly by the Universities of Edinburgh and Aberystwyth that led not only to new insights in robot control, but also stimulated novel research in related areas. The project was unusual for its time in that it integrated robotics expertise from a Computer Science department at Aberystwyth University in Wales with the interests of the Laboratory for Cognitive Neuroscience in the Psychology Department at Edinburgh University. Details of the project are given in the Acknowledgements.

Our main motivation was, and still is, to gain an understanding of how robot systems could achieve some of the rich, flexible behaviour seen everywhere in the autonomous agents of the animal kingdom. There still remains a large gulf between the behaviour produced by our best robotic efforts and the richness of behaviour, learning and adaptability so obviously manifest in living systems. In the 1990s we were unsatisfied with the current methods for designing and engineering of intelligent systems, and found a lack of general principles for embodied intelligence research. In particular we saw psychology as the potential missing link, with its higher-level

models, emphasis on behaviour and relative openness to the problems of complexity. Since then, we have developed our approach further and have followed three principled lines of research, each quite different but all relating in their own way to the central problem of understanding and designing complex autonomous systems.

In the next section, we briefly expand on the important issue of autonomy in robotics. Then follows a summary of our early founding work, before the three following sections each describe in turn our three lines of attack: steps towards a science of mobile robotics; an approach for developmental learning; and the potential role of full 3D geometric knowledge of the world.

2 Autonomy and Embodiment

The concept of Autonomous Systems can have many realisations but a central characteristic is the ability to sense, understand and act upon the environment in which the system operates. Thus, any internally processed information must be grounded in meanings ultimately derived by sensing and acting in the environment. This is why robotics is an excellent framework for autonomous systems research as it forces issues like sensing, perception, action, error-recovery and survivability, to be faced in an integrated and challenging format.

The important paradigm shift in robotics brought about by the Embodied movement has been the rejection of simplified “toy worlds” or artificial simulated environments and the emphasis on the “real world”. Furthermore, *unstructured* environments are the required proving ground for modern experiments, bringing not only realistic noise, disturbance and uncertainty to the fore, but also opening up the enormous complexity that autonomous systems must ultimately face and manage.

For example, consider the following scenario. A robot is to patrol the sea-facing pedestrian area at a popular coastal resort. The robot might be required to perform coastal surveillance, environmental audit, assist the public, monitor local conditions and search for missing persons (e.g. via heat sources). Such assistive robots will need a wide variety of sensors (to monitor local conditions; weather, waves, tide), some form of interface for interactions with the public, a link to remote services (control room or coastguard), and sophisticated perceptual processes (tracking and awareness functions to raise security warnings; novelty analysis for unusual events or objects).

It is clear that such a system must survive unattended for long periods thus requiring genuinely autonomous operation. But autonomy is not dependent on any group of specific functions or capabilities. Rather, autonomy is the ability to *cope* with changing situations and circumstances, and this in turn depends upon gaining a grounded understanding of those very situations from experience. For example, a sea-front robot that is temporarily unable to deliver a warning message directly may communicate with other robots or agents to recruit their help and achieve the task by different means. Similarly, actions that fail in one context (e.g. sandstorm) might be reconfigured through experience with other actions in related conditions (e.g. fog).

Embodiment is a vital property of autonomous systems because any understanding of the environment must be built up from sensory-motor activity and the morphology and physical hardware of a robot is an essential factor in determining both its behaviour and the extent of its cognitive abilities. We take the view of others [Lakoff, 1987] that all cognitive competencies are *grounded* in sensory-motor acts and even higher functions such as language are intimately related to basic sensory and motor experience [Rizzolatti and Arbib, 1998].

The achievement of autonomy is seen by the robotics community as one of the most pressing research challenges and is essential for the successful deployment of robots in service, domestic and health-care scenarios as well as in hostile and remote environments.

3 Control and Perception in Mobile Robots

In the early 1990s we began exploring new approaches for building autonomous robots working in everyday unstructured environments. Our joint task concerned the development of a courier robot capable of locating and identifying targets in order to carry out a given cognitive task, namely the delivery of mail or messages via mapping, route finding and navigation of its local environment.

We used mobile robots of various manufacture. Figure 1 shows a typical modern wheeled robot with a cylindrical body and a range of sensors (ultrasonic, infra-red, laser) mounted around its



Figure 1: A typical modern mobile

circumference. There are 2 main wheels, set on a diagonal, so that by driving these in contra or similar directions the robot can either rotate about its centre or move forward.

The difference between normal and unexpected events is of vital importance for autonomous systems and so one of our first investigations concerned the notion of novelty. We defined “novel” events as those not conforming to the current model of usually experienced events, and we built a variety of data-driven novelty filters that adapted through real-world robot-environment interaction. Using sonar as the sensing modality, a self-organising feature map performed clustering on the pre-processed sonar perceptions of the robot. This was demonstrated in a corridor exploring task where the robot initially only sees closed doors in a corridor and so treats open doors as “novel” but later treats both open and closed as “normal”.

This work was extended for dynamic environments where dynamical objects (e.g. other robots) are present and need to be learned as "normal". Using visual input from cameras on the robot we adopted a model-based approach which combines expectations from a model with current sensory perceptions. If either an abnormal sensory perception is perceived, or a behaviour is detected that is unusual in the robot's current context, the situation is classified as novel [Neto and Nehmzow, 2004] . A robot behavioural model was obtained using the Narmax system of non-linear polynomial identification (described in section 4.3) and then the robot was tested in a complex environment that had been encountered before. Novelty is signaled by a large error difference between the input perceptions and the model predictions. When any significant change was made to the environment, such as removing a small pillar, adding a barrier between two pillars or displacing a pillar, all were successfully detected as novel events [Ozbilge et al., 2009].

Development of Competencies		
BEHAVIOUR	MAPPING	LEVEL OF VISION
Reactive Stage		
go_forwards	none	threshold sonar
Locative Stage		
go_forwards	position of defaults	threshold sonar
go_forwards	sonar values + defaults	distance-value sonar
goto_point (using goForwards)	(uses map data)	distance-value sonar
Vision Stage		
findAny_object	object position + sonar values + defaults	object/non-object sensing + distance-value sonar
label_specific_object	object name and position + sonar values + defaults	distant object sensing + + feature identification + distance-value sonar
findspecific_object (uses map data) (using goto_point with avoid turned off)	as above	as above

Another emphasis was on the growth of competence necessary if an autonomous robot is to build on its experience. An illustration of this can be seen in our early implementation of a visual layer of competence. A camera and visual processing system on a mobile robot were used to

develop adaptive and learned responses to obstacles, recognise objects and recognise target locations [Nehmzow et al., 1993]. The cognitive levels in the final system consists of three major parts: a reactive layer, a locative competence, and a visual ability. The interaction, growth and development of these stages were our focus of interest. The three stages of competence can be summarised in the above table. The table shows the behaviours growing in complexity from a simple propensity to move forwards to a many-layered ability to find a specific object whilst avoiding intermediate objects and obstacles.

Sonar and vision competence also develops gradually from a simple threshold sonar device that can only ask the question 'am I too close to a surface?' to advanced vision that can discriminate between two different objects of interest to the agent and ask, 'is this the object I want?'

The mapping process largely reflects the maturation of the other two areas of development. It starts by plotting where defaults occur during the 'goForward' behaviour and then adds the information that it gets from the visual/sonar system as it becomes more complex. At the same time, the robot is using the lowest-level, and computationally simplest, visual information that it can to perform a task, and thus costly high level visual information is only used when necessary to open up new possibilities.

3.1 Control in organically unstructured environments

After taking account of relevant considerations from ethology we developed for robotics the servo-based model of behavioral control proposed by William Powers [Powers, 1973]. In classical control theory any feedback signals (usually negative) are fed back from the output of the controlled plant. Thus a motor speed controller will monitor the speed of the output shaft and use the error value to adjust the input power accordingly. The basic conceptual difference in Powers's model was to take the feedback signal *from the environment*, that is, *after* the effects of the plant had been felt on the environment. Thus the feedback loops are closed by environmental interaction. The difference from the behaviour-based methodology is that the feedback is now *within* the perceptual process, thus behaviour becomes the consequence of acting to reduce perceptual mismatch. We produced a variation on Powers theory, and implemented several models to investigate its properties [Rodrigues and Lee, 1994]. This approach offers a framework in which low level reactive behaviour can be integrated with higher level schemas in a control hierarchy - a challenging issue in Embodied Intelligence.




Figure 2: Sandwich assembly. A laser stripe (visible at extreme lower left) produces a profile of the product on an imaging camera when the belt is moving. This is then used to deduce the items that are missing and direct the robot to fetch and place such components of the sandwich.

We applied these ideas in various real applications with cluttered environments for both mobile robots and robotic manipulator arms and hands. One example was the development of robot grasping strategies for unprocessed natural food products, including packing fish portions into boxes and sandwich assembly. A typical task is to grasp and place natural food items at target locations with an industrial quality assembly robot, but with no prior knowledge of the product size

or shape. Natural food items are extremely variable in their shape, consistency and quality and these features preclude any preprogrammed grasping pattern. Using the Powers approach we designed a grasping algorithm that solves this problem by adapting the robot gripper orientation to perceived features of the product without operator intervention. Figure 2 shows a layer of cucumber and tomato slices being located during the assembly of a sandwich.

Another feature of this work was the application requirement for absolutely no programming to be allowed not only for the grasping of variable items but also for setting up the equipment for different batch runs. This was because of the hygiene regime and the management obsession to completely minimise human involvement. We achieved this by a system of "teaching by showing". A single exemplar sandwich would be shown to the camera system in each of its partial assembly stages, then the robot could deduce the action required by noting the difference between the current stage S_i and the next successive stage S_{i+1} . The robot action was automatically generated by the goal of reducing the difference to zero. In this way, any assembly sequence could be defined, and the quality of the result would reflect the quality of the examples initially shown to the system.

4 Scientific Methods in Robotics: Towards a theory of robot-environment interaction

The literature in robotics research contains many impressive feats of engineering. A typical paper will describe the problem analysis, design and performance results of a system that tackles some challenging task. The *modus operandi* has often been one of prototype building with successive refinement and the cumulative growth of expertise and knowledge (on the part of the experimenter). This process might characterise the very early days of Victorian engineering but is unsatisfactory for modern science. In particular, the most glaring omission is the lack of reproducibility of the published results. It is almost impossible to test claimed results by reproducing experiments because (a) the full details of the equipment, software and conditions are not given (usually because they would take far too much space), (b) it is impossible to duplicate the exact same laboratory apparatus as the equipment is often part original or modified, and (c) the initial conditions for experiments (including ambient conditions like lighting) are not fully recorded or are otherwise unavailable.

We believe that this situation must change if we are to attain a more scientific approach to building robot systems. We must build up an organised body of scientific knowledge that facilitates a much better understanding of such systems and allows for proper evaluation of our understanding and progress.

Taking mobile robotics as an example, assume that a small mobile robot is to be used in experiments to trace around the walls of an arena - this is known as a wall-following task. A program will cause the robot to act but the actual path taken will depend upon the combined influences of the control program running on the robot, the robot's physical properties, and the environment itself. Hence robot behaviour can not be reduced simply to the output of a program but is the result of interactions of the triple: Robot/Environment/Task. This is why robotics research is hard, as can be appreciated by the following case study based on real trials. Our wall-following robot is programmed to move parallel to any wall, while maintaining a constant distance from the wall. Unfortunately we soon find that the robot's path - its trajectory around the arena - is not repeatable. Each time we try to start it off in exactly the same location with exactly the same speed and heading we always find that after a while its path deviates from the previous one. Figure 3 shows an example of the deviation between two trajectories in four increasing time snapshots. At first the paths are near identical but they gradually diverge and eventually may become completely different trajectories. This effect is caused because of very slight variations in friction, material properties or ambient conditions. Such microscopic effects are unavoidable and uncontrollable, and consequently present a problem for our desire for scientifically repeatable results.

This problem is one of chaos — however careful we are, we have no way of controlling the conditions so that identical behaviour is produced — this is the sign of a chaotic system. Note that this is the case even though all the components of the experiment may be deterministic in nature. A solution might involve a way of measuring behaviour that can take account of these chaotic effects. We need a suitable "Behaviour Meter" that can be used to compare different behaviours. If we consider just the dynamics of our mobile robot behaviour then various possible quantitative methods can be considered. We have experimented with quantitative descriptions of phase space

and found three particular measures very helpful [Nehmzow, 2009].




Figure 3: Four snapshots of two trajectories. Starting from top left, time increases clockwise

4.1 Phase space

Phase space is a concept from dynamical systems theory and is used to describe all possible states of a dynamical system. Not only spatial position but velocities will typically be involved, thus phase spaces may have high dimensionality. It turns out that if a log is taken of robot variables at regular time intervals during a behavioural trajectory then the phase space for the system can be reconstructed from the observations in these timeseries alone. Figure 4 shows an example of part of a reconstructed phase space (this is a 5 dimensional phase space but only 3 can be shown). Note that this is not a robot spatial trajectory but a plot in phase space — the variation in the traces indicates the degree of unpredictability; a strictly periodic behaviour would produce a single curve.

4.2 A behaviour meter

Given a phase space we can then apply our three measures from our "Behaviour Meter"; these are, the Lyapunov exponent, the Prediction Horizon, and the Correlation Dimension. The Lyapunov exponent is a measure of the phase space that describes the rate of divergence of two trajectories that started infinitesimally close to each other. Figure 5 shows the trajectories of a mobile robot executing two different behaviours: wall-following and obstacle-avoidance. It is clear that wall-following is the more predictable behaviour as the trace shows. Obstacle-avoidance causes the wall to be avoided as soon as it is detected and therefore looks more unpredictable and chaotic. The Lyapunov exponent for the wall-following case was calculated as between 0.02 and 0.03 bits/sec, while in the obstacle-avoidance trial it was between 0.11 and 0.13 bits/sec. This shows that obstacle-avoidance is indeed more chaotic than wall-following.

The Lyapunov exponent is expressed as information loss per unit time; in other words it indicates the loss of information in the system (or degree of chaos) as one predicts the system state for longer and longer times ahead. We may be able to predict our robot's state in the next second fairly accurately but we would expect large errors if we try to predict several hours ahead. The Lyapunov exponent tells us how bad the situation is — if the Lyapunov exponent is zero then we have a noise-free, perfectly deterministic system whose behaviour we can accurately predict for any length of time. But as the Lyapunov exponent increases so our predictions get worse and eventually become no better than a guess. We refer to the point in time where complete loss of predictability occurs as the Prediction Horizon. Using information theory analysis we find that the Prediction Horizon for the wall-following behaviour was greater than 25 minutes (6000 steps at sampling rate of 4Hz) and for the obstacle-avoidance was only 80 seconds. This shows that wall-following is much more predictable and for any model of obstacle-avoidance, however good the model, it will not be able to predict the exact path of the robot for more than about 80 seconds.

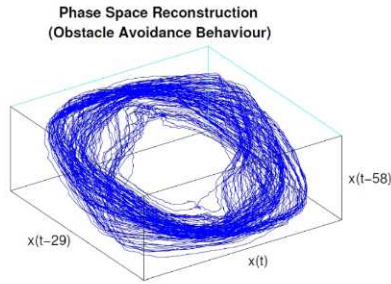


Figure 4: Phase space example (showing 3 of 5 dimensions)

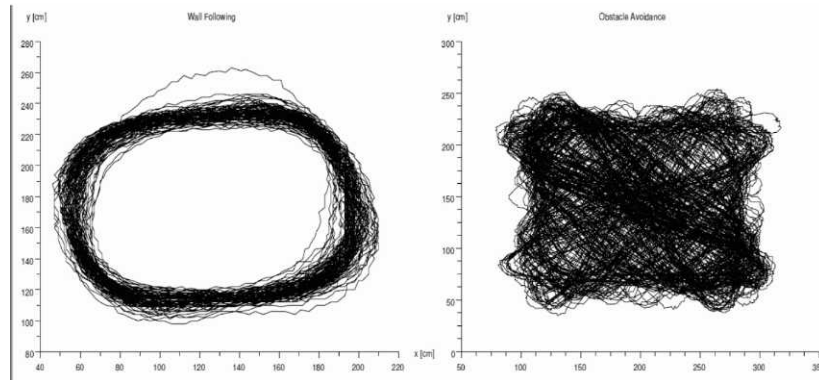


Figure 5: Trajectories of wall-following and obstacle-avoidance as recorded from an overhead camera.

The final measure of chaos is the Correlation Dimension, which gives a measure of aperiodicity or how close the system variables return to previous values (zero indicates strictly periodic data). For the wall-following behaviour the Correlation Dimension was calculated as around 1.5 while for obstacle-avoidance it was 2.5, again clearly indicating the increased chaotic nature of the obstacle-avoidance behaviour [Nehmzow, 2009].

4.3 Faithful and transparent modelling of robot-environment interaction

The issues we have discussed regarding real robot behaviour have some important implications for modelling and simulation. For a model to have any value it must be accurate and faithful, that is, it must predict or generate identical behaviour to the originally perceived behaviour. Models should also be reasonably transparent and analysable so that we can understand their meaning and what they represent. As an example of a non-transparent model consider the case of a neural network that has been well trained on some specific task. The only information we have available is the weight matrix that has captured the essence of the input-output relationship. But unfortunately the weight matrix is just an array of numbers that obscurely encode whatever the model has learned; they are quite unintelligible to humans. By comparison we are interested in transparent, parsimonious models that aid understanding of these complex situations.

Clearly, if there are indications of chaotic behaviour then the predictability of any models may be severely limited. However it is possible to discern structure in even noisy data in many cases. The field involved in producing models of unknown systems is called System Identification, and a powerful and relevant method for robotics is the Narmax technique [Chen and Billings, 1989]. This is a mathematical method that produces a non-linear polynomial of the system variables that expresses the relationship between a set of inputs and outputs. As a simple illustration consider a system with two inputs, u_1 and u_2 , and a single output, y . We chose the system equation as:

$y(n) = \mathbf{u}_1(n)^2 + 0.5\mathbf{u}_1(n)\mathbf{u}_2(n)$, where n is the timeseries index. Thus this gives the ideal or theoretical output for any input values. But in the real life situation, (a) the inputs contain a lot of

noise and (b) we have no idea of any equation that might represent the system. This is where system identification methods are useful and to demonstrate we use the Narmax method [Chen and Billings, 1989]. First we collect some data by generating a set of values for u_1 and u_2 and add Gaussian distributed noise giving u_1 and u_2 , then we pass this data through our system (the above equation) to produce the output data y . We now have a set of input-output data for a noisy system. This data is then processed through the Narmax method, which delivers the following result:

$$y(n) = -0.01 + 0.05u_1(n) + 0.01u_2(n) + 0.91u_1(n)^2 + 0.52u_1(n)u_2(n)$$

This is very close to the true underlying relationship: $y(n) = u_1(n)^2 + 0.5u_1(n)u_2(n)$, and we now have a model of the system for further analysis, simulation or prediction [Nehmzow, 2009].

4.4 Robot control

We have applied these techniques to a number of standard robotic problems. An interesting approach is to build the system model from data collected while human operators drive a robot to achieve a particular task. The resulting model then contains the expertise of the operators and can be used as a controller to drive the robot autonomously. Figure 6 shows a doorway navigation task viewed from above. On the left an operator has driven the robot from many starting points (in the region at the top) and navigated through the gap in the wall (the gap is 2 robot diameters wide). Narmax models produce polynomials with many terms, for the various combinations of inputs and their products. In this case there were many laser and sonar inputs and the model polynomial came out with 38 terms. The robot was then run under control of the model, taking its actions from the model processed sensory input. The traces in Figure 6 show 39 runs and those of the human (on the left) are noticeably less smooth than those produced by the model (on the right). Doorway traversal is a delicate task, requiring a careful balance between several signals, but the Narmax model was successful every time. Post analysis of the model showed that just a few of the sensors were playing a major role and these were all monitoring one side of the doorway. This strategy of following one side closely is clearly seen in Figure 6 and was an unexpected effect. The model driven robot was much smoother and more accurate than a human operator!

5 The Importance of Development for Cognitive Robots

We have argued that true autonomy involves dealing with the new or unknown without external aid. This means that systems must not only adapt in accordance with current experience but must also be capable of adapting their learning processes themselves. Thus, new competencies must emerge as conditions change and new demands are made. The way this problem has been solved in



Figure 6: Robot under manual control..... and under model

humans and

other mammals is through processes of structured growth generally known as "development". In this section we consider the role of development in relation to cognitive aspects of robotics.

Developmental psychology has long studied human cognitive growth and produced many

theories that could explain such growth. It is very surprising that, despite the vast body of psychological knowledge on learning and adaptation built up over the last century, very little work has considered implementing developmental processes in artificial systems.

This situation has finally changed and the topic of Developmental Robotics has recently become established as a new research area [Lungarella et al., 2003]. This approach emphasises the role of environmental and internal factors in shaping adaptation and behaviour, and posits a developmental framework that allows the gradual consolidation of control, coordination and competence [Prince et al., 2005].

5.1 Developmental stages

A key characteristic of human development is the centrality of behavioural sequences: no matter how individuals vary, all infants pass through sequences of development where some competencies always precede others. This is seen most strongly in early infancy as one pattern of behaviour merges into another. These regularities are the basis of the concept of behavioural *stages* — identifiable periods of growth and consolidation. Perhaps the most influential theories of staged growth have been those of Jean Piaget who emphasised the importance of sensory-motor interaction, staged competence learning and a constructivist approach [Piaget, 1973]. It is recognised that stages tend to have vague boundaries and also vary greatly with individuals. Nevertheless, the existence of stages in development and their role in the growth of cognition appears to be very significant.

We believe that research into developmental algorithms for robotics should begin with and be firmly rooted in the early sensory-motor period. This is for several reasons: (1) it is logical and methodologically sound to begin at the earliest stages because early experiences and structures are highly likely to determine the path and form of subsequent growth in ways that may be crucial; (2) according to Piaget, the sensory-motor period consists of six stages that include concepts such as motor effects, object permanence, causality, imitation, and play — these are all issues of much relevance to robotics; (3) sensory-motor adaptation and learning is vital for autonomous robots; (4) it seems likely that sensory-motor coordination is a significant general principle of cognition [Pfeifer and Scheier, 1997].

Hence, we are investigating the earliest level of sensory-motor development: the emerging control of the limbs and eyes during the first three months of life. To the casual observer the newborn human infant may seem helpless and slow to change but, in fact, this is a period of the most rapid and profound growth and adaptation. From spontaneous, uncoordinated, apparently random movements of the limbs the infant gradually gains control of the parameters, and learns to coordinate sensory and motor signals to produce purposive acts in egocentric space [Gallahue, 1982]. We believe there is much to learn for Embodied Intelligence and robotics from this scenario.

5.2 The key role of constraints

Any constraint on sensing, action or cognition effectively reduces the complexity of the inputs and/or possible action. This reduces the task space and provides a frame or scaffold which shapes learning [Bruner, 1990, Rutkowska, 1994]. When a high level of competence at some task has been reached then a new level of task or difficulty may be exposed by the lifting of a constraint [Rutkowska, 1994]. The next stage then discovers the properties of the newly scoped task and learns further competence by building on the accumulated experience of the levels before.

Various examples of internal sensory and motor constraints are seen in the newborn, for example the neonate has a very restricted visual system, with a kind of tunnel vision [Hainline, 1998] where the width of view grows from 30 degrees at 2 weeks of age to 60 degrees at 10 weeks [Tronick, 1972]. Although this may seem restricted, these initial constraints on focus and visual range are "tuned" to just that region of space where the mother has the maximum chance of being seen by the newborn. When "mother detection" has been established then the constraint can be lifted and attention allowed to find other visual stimuli.

Many forms of constraint have been observed or postulated [Hendriks-Jensen, 1996] [Keil, 1990] and we have identified a range of different types in robotics. These include: anatomical or hardware constraints imposed by the system morphology; sensory-motor limitations (e.g. accuracy,

resolution, bandwidth); cognitive/computational constraints; maturational constraints from internal and biological processes; and, not least, external or environmental constraints.

In our current research we are following these ideas of staged development and building robotic learning architectures in which sensory-motor competence grows cumulatively. Our method involves an overarching constraint network that restricts the ranges and number of parameters available to the robot during its early stages.

5.3 The LCAS approach

As mentioned before, we try to find explicit, abstract models and avoid preselected internal representations and methods. Accordingly, we have produced a general mechanism for staged development which takes “constraint lifting” as a key process in allowing transitions between stages. We use novelty and expectation as the drivers and so any trigger for stage transitions is likely to be related to internal global states, not local events. Local stimuli (spatially and temporally) may cause local responses but global values can indicate levels of general experience or expectation.

Thus, global states such as global excitation can act as indicators that can detect qualitative aspects of behaviour such as when growth changes have effectively ceased or when a mapping between modalities has become saturated. They can then signal the need to enter a new level of learning by lifting a constraint or accessing a new sensory input. In this way, further exploration may begin for another skill level, thus approximating a form of Piagetian learning.

Our approach then consists of implementing the cycle; Lift-Constraint, Act, Saturate (LCAS), at a suitable level of behaviour. First, the possible or available constraints must be identified and a schedule or ordering for their removal decided. Next a range of primitive actions must be determined together with their sensory associations. Also any sensory-motor learning or adaptation mechanism is incorporated at this stage. Finally a set of global measures need to be established to monitor internal activity. When this is implemented the initial behaviour may seem very primitive, but this is because all or nearly all constraints have been applied and there is little room for complex activity. During the Act process varying patterns of action effectively explore the scope for experience and any new experiences are learned and consolidated. Eventually there are no new experiences possible, or they are extremely rare, and this level becomes saturated. The global indicators then reach a critical level and the next constraint in the schedule is lifted and the cycle begins again.

The ideas reported here have all been explored in experiments on hand/eye robot systems. The details cover the learning of sensory-motor control for eye-saccades [Chao et al., 2010], visual search [Huelse et al., 2009a], and hand-eye coordination [Huelse et al., 2009b, Huelse et al., 2010]. The behaviours observed from our experiments display an increasing progression from initially spontaneous limb movements (known as "motor babbling"), followed by more exploratory movements, and then directed action towards touching and grasping objects. Our research is continuing in a programme that aims to demonstrate autonomous cognitive growth on an iCub humanoid robot [Metta et al., 2008]. We

are exploring constraint networks as an overarching framework for orchestrating development and have build such networks by transposing a large sensory-motor constraint analysis of the human infant.

6 The Role of Geometric Knowledge in Recognition Tasks

The application described in section 3.1, involving robotic handling of natural food products, operated very successfully with only 2 dimensional images of the environment. No prior information about the product shape was available and the only knowledge used was a set of control variables in a geometric relationship and the sensory input.

A follow-on research question addresses the opposite end of the knowledge spectrum: would full geometric knowledge of the shape allow us to explore in more detail the perception/anticipation/action model? In order to address this question, objects in the real world need to be directly perceived by artificial sensors and their geometries reconstructed from such perceptions.

Our research has focused on 3D data acquisition and exploitation from single 2D images using structured light methods, e.g. [Robinson et al., 2004, Brink et al., 2008, Rodrigues and Robinson,

2009]. We demonstrate that a robot can effectively track targets in 2D and reconstruct these into 3D as it wanders through the environment.

6.1 2D tracking and 3D reconstruction

The OpenCV Intel libraries [Bradski and Pisarevsky, 1999] provide built-in routines for real-time face detection based on Haar-like features. It is possible to train and use a cascade of boosted classifiers for rapid object detection for any arbitrary object. We have used the libraries to train classifiers for left and right eye. The general problem with such detection techniques is the number of false positives. For instance, in any image there could be various detected faces and some might not be real faces. Similarly for eyes, the routines normally detect more eyes than there are in the scene.

The solution is to run face and eye detection in separate threads and impose constraints: first, there should be only one face detected in the image and the face must be larger than a certain threshold; second, there should be only one left and only one right eye detected in the image, and these must be within the region of interest set by the face detection; third, the position of the face and eyes must not have moved more than a set distance since last detection so to avoid taking blurred shots due to undesirable motion. The routines are thus dedicated and continuously track multiple face and eyes. Only when the above constraints are satisfied a shot is automatically taken; an example is depicted in Figure 7 (left).

We have developed a suite of routines for real-time 3D reconstruction using structured light patterns [Robinson et al., 2004]. To avoid the need for accurate mechanisms and in order to speed up the acquisition process, a number of stripes can be projected at the same time and captured as a sequence of stripes in a single frame. However, it may be difficult to determine which captured stripe corresponds to which projected stripe, when we attempt to index the captured sequence in the same order as the projected sequence. We call this the stripe indexing problem. For this reason methods have been devised to uniquely mark each stripe, by colour [Rocchini et al., 2001], stripe width [Daley and Hassebrook, 1998], and by a combination of both [Zhang et al., 2002].

Our research has shown the dependence between the stripe index and the measurement of a surface vertex defined by the common inclination constraint [Robinson et al., 2004]. We also deal with occlusions [Wang and Oliveira, 2007] to improve the validity of the boundaries. Moreover, we have investigated how far the indexing problem can be solved with uncoded stripes, where the correct order of stripes in the captured image is determined by original algorithmic methods such as the maximum spanning tree algorithm [Brink et al., 2008]. A number of 3D post-processing techniques can be applied such as the ones discussed in [Rodrigues and Robinson, 2009, Wang and Oliveira, 2007] resulting in 3D models as depicted in Figure 7 (right).

6.2 Pattern recognition

Object recognition is based on feature extraction starting from three key feature points: the location of the eyes and tip of the nose. Our method is based on cutting oriented planes from the key features and detecting points on the mesh at the interception of those planes. We require pose alignment where the origin is placed at the tip of the nose, the eyes are aligned with the x-axis, and the y-axis is at a constant angle to a point at mid-distance between the eyes. This is achieved through an automatic iterative process, which has proved to work successfully even if the subject is not directly facing the camera; it has been tested on images facing up to 45 degrees to either side. A total of 43 points are located at the intersection of the various planes defined from the key features. An example of such points is depicted in Figure 8. Measurements are taken from such points as distances and ratios in addition to area, volume, perimeter, and various types of diameters such as breath and length resulting in a set of 191 measurements per face model. The face models have been tested and recognized at 97% accuracy for a database containing 276 models.




Figure 7: Left: when 2D constraints are satisfied for face and eye tracking, structured light is projected. Right: the captured 2D image is processed into 3D and texture mapping can be changed by 3D postprocessing operations.




Figure 8: Automatic pose alignment and feature

6.3 Discussion

Our research has demonstrated that defining specific robotic recognition tasks can be achieved in two distinct stages using appropriate constraints: 2D is well-suited for tracking objects in real time based on redundancy of information. Many objects (e.g. faces and eyes) can be tracked from a 2D scene and the actual selection of a particular scene for 3D reconstruction is made using constraints of size and number of objects detected.

Once a scene is reconstructed in 3D, then more specific knowledge is required about the object of interest. We used the example of facial recognition that relies on geometric knowledge of a face object for pose normalization and feature extraction. The need for further knowledge is likely to be true for the recognition of most 3D objects as their boundary conditions are only partially controlled through 2D tracking.

While the original question on deploying the perception/anticipation/action model remains largely unanswered, our research into 3D has opened new avenues of investigation. This includes games, animation, entertainment, security and engineering among others. The limitations of the method are related to projection issues, as objects over 4m from the camera cannot be reconstructed. We are currently working on new techniques to project sharp stripes over longer distances.

7 Summary

In our research journey we have learned much about the nature of autonomy and embodiment. It is interesting that our original research themes have continued, albeit in much advanced form, but still focussed on certain key themes. We summarise these briefly.

In the quest for more fluid, flexible and animate behaviour from our robots, we need to continue with the Embodied Intelligence approach to building artificial cognitive systems because the nature of the physical hardware of a robot is an essential factor in determining both its behaviour and the extent of its cognitive abilities. In addition, Developmental Robotics starts from the assumption that early experience provides a vital grounding for later competencies and thus offers an approach into the difficult problem of the growth of competence. Constraint networks could provide an environment in which behaviour develops without complex mechanisms for each stage, in other words: "Gradual removal of constraint could account for qualitative change in behaviour without structural change" [Tronick, 1972].

Our interest in novelty as a learning spur has been justified by much current interest in Intrinsic Motivation [Oudeyer et al., 2007] which is examining the drivers for self-motivation, of which expectation and novelty are integral aspects.

We have seen the value of insisting on natural, unstructured environments, and the high degree of complexity that this may entail. We have argued that complexity and chaos need to be faced, not avoided, and that tools for measuring behaviour in such situations are essential for good science. Our examples included tools to measure the degree of chaos and sensitivity to initial conditions. We hope that more robotic science and high quality results will emerge from the use of new measurement techniques, proven authentic models, and take maximum advantage of psychological knowledge of cognitive growth.

8 Acknowledgements

The early joint work reported in section 3 was carried out between 1993 and 1996 and was supported by the MRC through grant G9110835/G9125255 entitled "Cognitive Modelling and the Design of Artificially Intelligent Systems". The work was undertaken in the Centre for Intelligent Systems at the University of Wales, Aberystwyth (Principal investigator: Lee) and in the Department of Psychology at the University of Edinburgh (Principal investigator: Dr Brendan McGonigle). We also gratefully acknowledge the BBSRC, EPSRC and the European Commission for their support of many subsequent projects which developed the research further as described in sections 4, 5 and 6. We also especially appreciated the help, advice and encouragement of very many colleagues, students and friends.

References

- [Bradski and Pisarevsky, 1999] Bradski, G. and Pisarevsky, V. (1999). Intel's computer vision library: applications in calibration, stereo, segmentation, tracking, gesture, face and object recognition. In *IEEE Computer Science Conference on Computer Vision and Pattern Recognition*, volume 2. IEEE Computer Society; 1999.
- [Brink et al., 2008] Brink, W., Robinson, A., and Rodrigues, M. (2008). Indexing Uncoded Stripe Patterns in Structured Light Systems by Maximum Spanning Trees. In *British Machine Vision Conference*, pages 575-584.
- [Bruner, 1990] Bruner, J. (1990). *Acts of Meaning*. Harvard University Press, Cambridge, MA.
- [Chao et al., 2010] Chao, F., Lee, M., and Lee, J. (2010). A developmental algorithm for ocular-motor coordination. *Robotics and Autonomous Systems*, 58:239-248.
- [Chen and Billings, 1989] Chen, S. and Billings, S. A. (1989). Representations of non-linear systems: The narmax model. *Int. J. Control*, 49:1013-1032.
- [Daley and Hassebrook, 1998] Daley, R. and Hassebrook, L. (1998). Channel capacity model of binary encoded structured light-stripe illumination. *Applied Optics*, 37:17.
- [Gallahue, 1982] Gallahue, D. (1982). *Understanding Motor Development in Children*. John Wiley, NY.
- [Hainline, 1998] Hainline, L. (1998). How the visual system develops. In Slater, A., (Ed.), *Perceptual development: Visual, auditory, and speech perception in infancy*, pages 5-50. Psychology Press.
- [Hendriks-Jensen, 1996] Hendriks-Jensen, H. (1996). *Catching Ourselves in the Act*. MIT Press, Cambridge, MA.
- [Huele et al., 2009a] Huele, M., McBride, S., and Lee, M. (2009a). Implementing inhibition of return: embodied visual memory for robotic systems. In *Proc. 9th Int. Conf. on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems. Lund Uni. Cognitive Studies*, volume 146, pages 213-214.
- [Huele et al., 2009b] Huele, M., McBride, S., and Lee, M. (2009b). Robotic hand-eye coordination without global reference: A biologically inspired learning scheme. In *Proc. Int. Conf. on Developmental Learning, ICDL*.
- [Huele et al., 2010] Huele, M., McBride, S., and M., L. (2010). Fast learning mapping schemes for robotic handeye coordination. *Cognitive Computation*, 2(1):1-16.
- [Keil, 1990] Keil, F. (1990). Constraints on constraints: Surveying the epigenetic landscape. *Cognitive Science*, 14(4):135-168.
- [Lakoff, 1987] Lakoff, G. (1987). *Women, Fire, and Dangerous Things*. University of Chicago Press, Chicago, USA.
- [Lungarella et al., 2003] Lungarella, M., Metta, G., Pfeifer, R., and Sandini, G. (2003). Developmental robotics: a survey. *Connection Science*, 15(4):151-190.
- [Metta et al., 2008] Metta, G., Sandini, G., Vernon, D., Natale, L., and Nori, F. (2008). The iCub humanoid robot: an open platform for research in embodied cognition. In *PerMIS: Performance Metrics for Intelligent Systems Workshop. Aug*, pages 19-21.
- [Nehmzow, 2003] Nehmzow, U. (2003). *Mobile Robotics: A Practical Introduction*. Springer Verlag.
- [Nehmzow, 2009] Nehmzow, U. (2009). *Robot behaviour: design, description, analysis and modelling*. Springer.
- [Nehmzow et al., 1993] Nehmzow, U., Smithers, T., and McGonigle, B. (1993). Increasing behavioural repertoire in a mobile robot. In Meyer, J., Roitblat, H., and Wilson, S., (Eds.), *From Animals to Animats 2*, pages 291-297. MIT Press.
- [Neto and Nehmzow, 2004] Neto, H. and Nehmzow, U. (2004). Visual novelty detection for inspection tasks using mobile robots. In *Proc. Brazilian Symposium on Neural Networks, IEEE Computer Society, Los Alamitos*.
- [Newell and Simon, 1976] Newell, A. and Simon, H. (1976). Computer science as empirical enquiry: Symbols and search. *Communications of the ACM*, 19:113-126.
- [Oudeyer et al., 2007] Oudeyer, P., Kaplan, F., and Hafner, V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation*, 11(2):265-286.
- [Ozbilge et al., 2009] Ozbilge, E., Nehmzow, U., and Condell, J. (2009). Expectation-based novelty detection in mobile robotics. In *Proc "Towards Autonomous Robotic Systems"*

(TAROS).

- [Pfeifer and Scheier, 1997] Pfeifer, R. and Scheier, C. (1997). Sensory-motor coordination: the metaphor and beyond. *Robotics and Autonomous Systems*, 20(2):157-178.
- [Piaget, 1973] Piaget, J. (1973). *The Child's Conception of the World*. Paladin, London.
- [Powers, 1973] Powers, W. (1973). *Behavior: The control of perception*. Aldine Chicago.
- [Prince et al., 2005] Prince, C., Helder, N., and Hollich, G. (2005). Ongoing Emergence: A Core Concept in Epigenetic Robotics. In *Proceedings of the 5th Int. Workshop on Epigenetic Robotics*, pages 63-70. Lund University cognitive studies.
- [Rizzolatti and Arbib, 1998] Rizzolatti, G. and Arbib, M. (1998). Language within our grasp. *Trends in neurosciences*, 21(5):188-194.
- [Robinson et al., 2004] Robinson, A., Alboul, L., and Rodrigues, M. (2004). Methods for indexing stripes in uncoded structured light scanning systems. *Journal of WSCG*, 12(3):371-378.
- [Rocchini et al., 2001] Rocchini, C., Cignoni, P., Montani, C., Pingi, P., and Scopigno, R. (2001). A low cost 3D scanner based on structured light. In *Computer Graphics Forum*, volume 20, pages 299-308. Eurographics.
- [Rodrigues and Robinson, 2009] Rodrigues, M. and Robinson, A. (24-27 July 2009). Novel methods for real-time 3D facial recognition. In *5th Int Conf on Comp Science and Info Systems*.
- [Rodrigues and Lee, 1994] Rodrigues, M. A. and Lee, M. H., (Eds.) (1994). *Perceptual Control Theory. Proceedings of the 1st European Workshop on Perceptual Control Theory*. The University of Wales, Aberystwyth, Wales.
- [Rutkowska, 1994] Rutkowska, J. (1994). Scaling up sensorimotor systems: Constraints from human infancy. *Adaptive Behaviour*, 2:349-373.
- [Tronick, 1972] Tronick, E. (1972). Stimulus control and the growth of the infant's effective visual field. *Perception and Psychophysics*, 11(5):373 - 376.
- [Wang and Oliveira, 2007] Wang, J. and Oliveira, M. (2007). Filling holes on locally smooth surfaces reconstructed from point clouds. *Image and Vision Computing*, 25(1):103-113.
- [Zhang et al., 2002] Zhang, L., Curless, B., and Seitz, S. (2002). Rapid shape acquisition using color structured light and multi-pass dynamic programming. In *The 1st IEEE International Symposium on 3D Data Processing, Visualization, and Transmission*, pages 24-36.