The 'Natural Selection' of Robotics

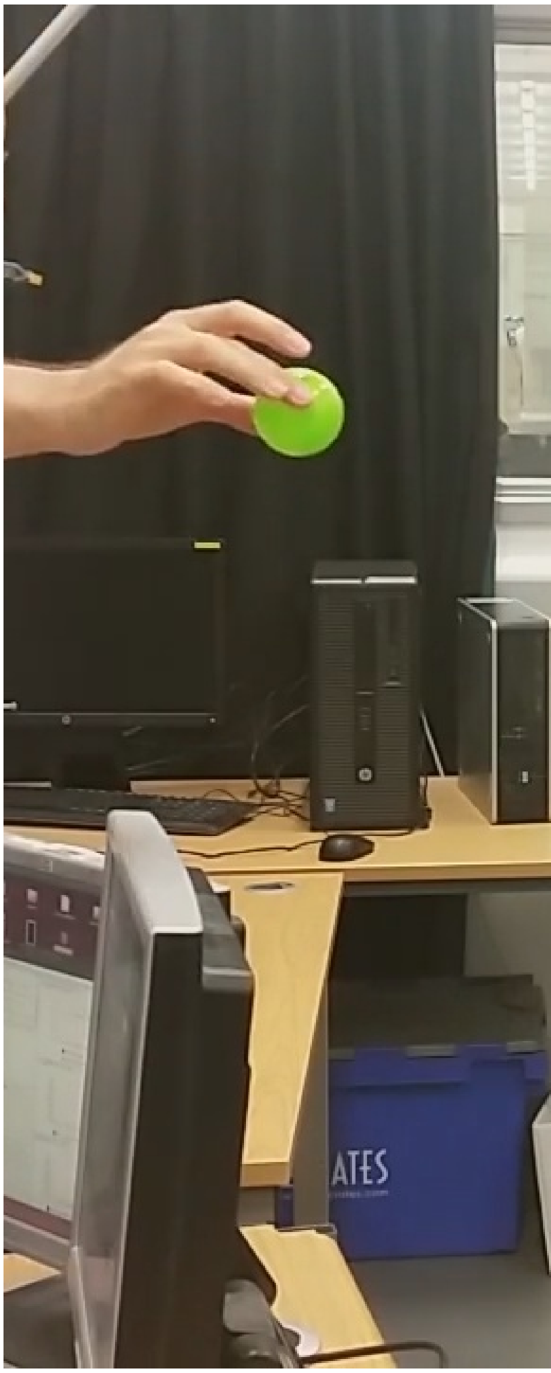

A fundamental revision of the robotics paradigm.

Rupert Young *reports.*

Prior to the middle of the 19th century the field of biology was in disarray. There was no coherent explanation of the diversity seen in the animal kingdom and the underlying biological processes from which that diversity evolved. Robotics and Artificial Intelligence (AI) is in a similar pickle today. Throughout the history of AI there have been various conceptualisations of what constitutes intelligent systems; knowledge-based systems, pattern recognition, behaviour-based robotics, neural networks, model-based control. However, there has been no theory that provides a language and architecture that unifies the different perspectives of intelligent systems. Despite all the hype around an imminent, mystical singularity practical applications of AI systems are unable to cope with the uncertainties and complexity of the real world (see **DARPA**) and remain confined to restricted and highly controlled environments.

Charles Darwin fundamentally revolutionised the entire understanding of the processes that led to the enormous diversity of life. His concept of natural selection provided a language by which evolutionary changes could be explained and understood, and an architecture of the underlying processes responsible. This simple, elegant concept demolished the prevailing supernatural view of the origins of life. As the geneticist Theodosius Dobzhansky said, "Nothing in biology makes sense except in the light of evolution".

"Nothing in biology makes sense except in the light of evolution."

A similar revolution is due in the field of Robotics and Autonomous Systems with the concept of *perceptual control*. Originally devised by the American physicist and engineer William T. Powers, **Perceptual Control Theory** (PCT) has common historical roots with cybernetics, but resolves "Weiner's Error" of the misapplication of feedback control to living systems. The standard control system model (see **Cybernetic Control** (a)), employed by Weiner, indicates that it is output that is controlled. From an engineer's point of view that is correct, it is what they observe as the response of the system; speed of car, for example.

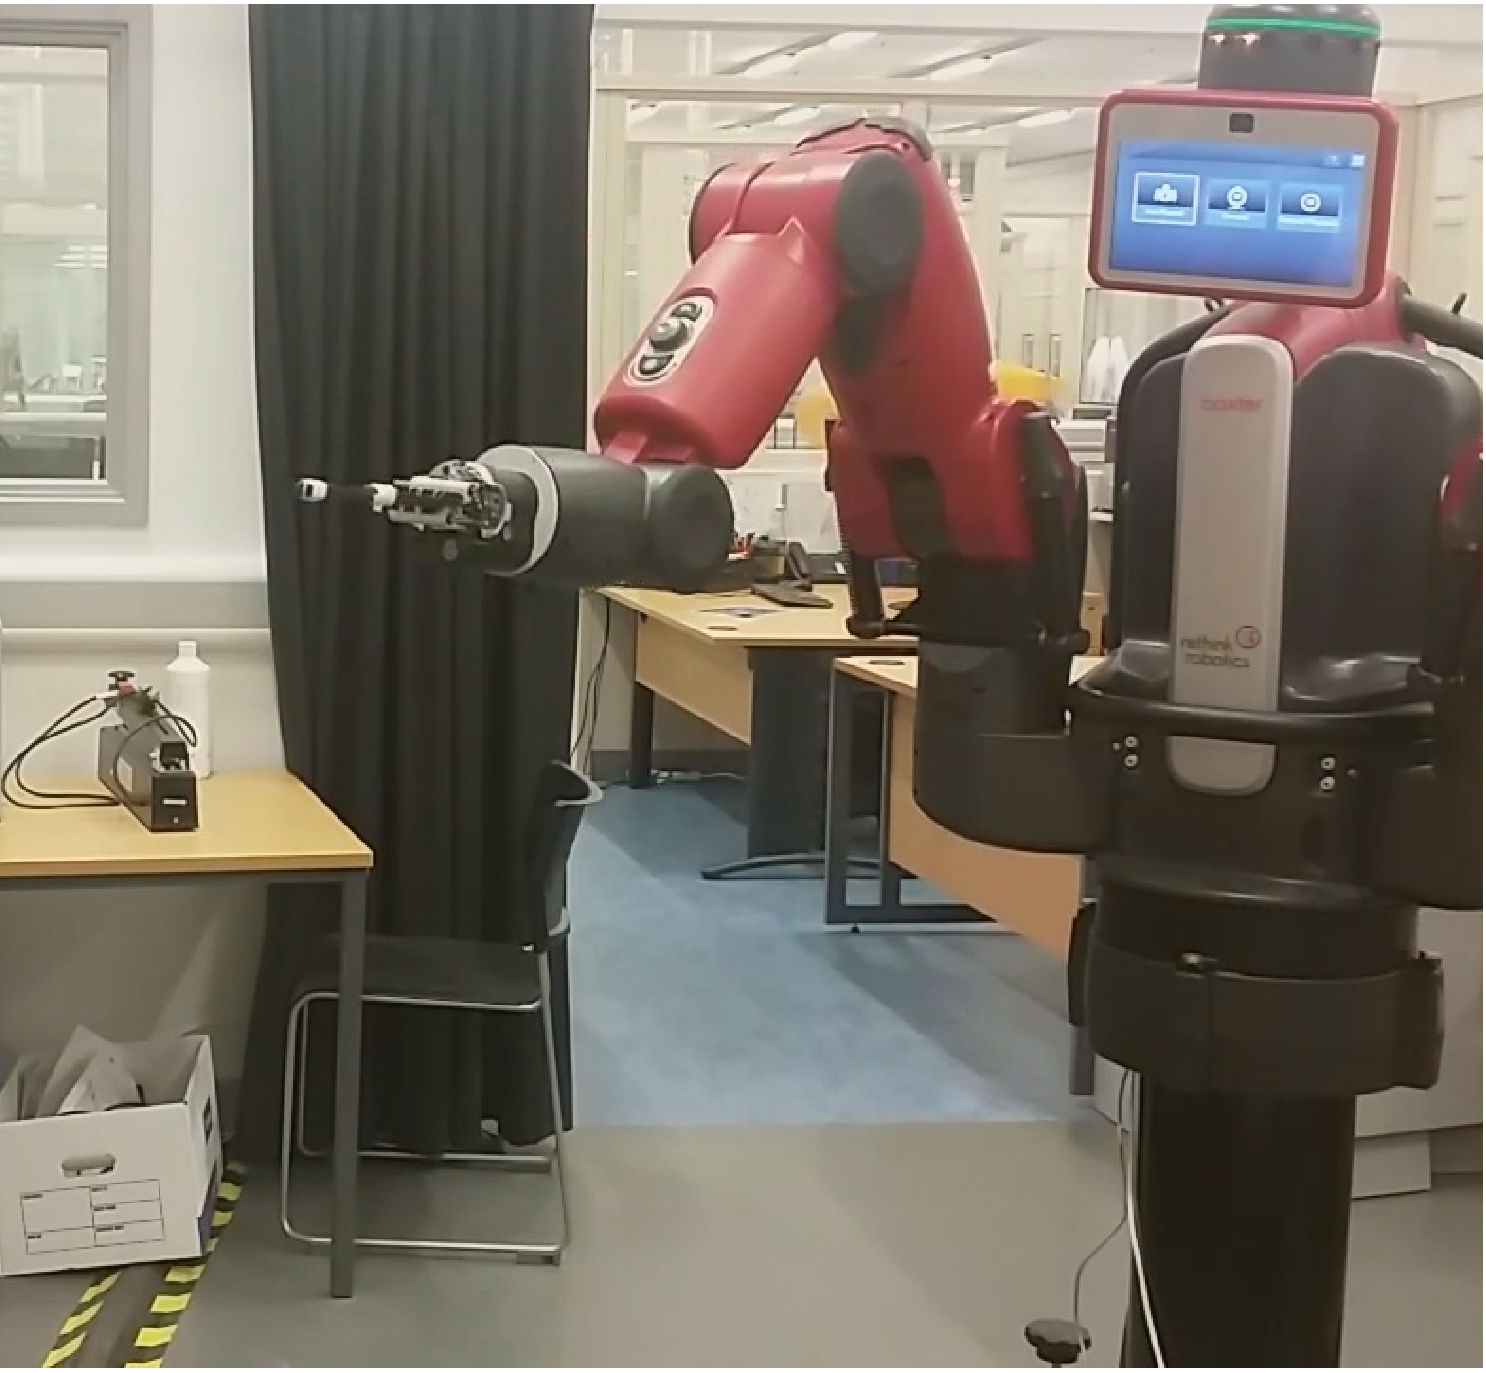

However, when applied to living systems output has been incorrectly interpreted as action. If we look again at the standard control model we can see that the branch from output to the comparator (see **Cybernetic Control** (b)) is actually the measured (or sensed) value of the variable under control. In other words, in terms of living systems the variable that is under control is the perception not action.

This misconception - that in order to operate, artificial systems need to control (compute) their behavioural output (action) - is not just a problem within Cybernetics.

The prevailing view within Robotics and AI goes way back to Alan Turing's computational approach whereby a future state of a system is computed from a current state via an *internalised model* of the external world. This has been highly successful for the predictable, highly deterministic environments of computers and simulations, but the extension of the metaphor to artificial systems operating autonomously in the real, dynamic, chaotic world has been an abject failure. Any progress that has been made over the last few decades has been more to do with advances in processing power than with an understanding of the underlying concepts of behaviour within living and intelligent systems.

In my paper, **A General Architecture for Robotics Systems: A Perception-Based Approach to Artificial Life**, published in the Artificial Life journal in Spring 2017, I describe how the concept and architecture of perceptual control can be applied to robotic systems drastically overhauling the methodologies currently in vogue.

In contrast to the standard behavioural view of computing output (action), perceptual control turns convention on its head with the central concept of controlling and maintaining a desired perception, by varying output. For example, when driving, our goal is to perceive the car between the white lines. We turn the steering wheel until we perceive that goal. We don't turn the wheel a specific amount or to a specific angle, but we keep turning the wheel to reach and maintain our perceptual goal. In other words, the goal state is a perception.

However, how we reach that state is not something we can predict in advance as there are many external, unknown factors that can affect that goal state, such as tyre pressures, road surface and wind. But we do not need to know the effects of those disturbances as we can indirectly detect their combined impact on the goal perception of the car's position and we *act* to protect that perceptual variable from any disturbance.

This highlights a fundamental flaw in the conventional model-based approach to intelligent systems. In order to compute the next state of the system it would be necessary to model everything that could possibly have an effect on the system. That is fine in the restricted domain of a game of chess or computer simulation, or even the highly controlled conditions of the factory floor, but with the real world you would have to be a **Laplacian Demon** (see panel **Simulation Problem**) that knew the entire state and dynamics of the universe.

This issue has not been recognised due to the **Simulation Problem**. In a simulation (usually the first stage of robot development) the relationship between the steering wheel and the heading can easily be defined and a simulated car can be controlled perfectly by this approach. But that is because with simulations we are acting as if we were Laplacian Demons, as it us who are defining the universe in which the simulation is running.

The situation is similar with controlled environments such as the factory floor or the laboratory, as the uncertainty is limited and managed, allowing us to define some, relatively modest, models. However, in the real world we can't do this; we can't be Laplacian Demons, and there are no such relationships to model. This is why robots have mostly been stuck in predictable, structured environments and why the modelling approach is not viable in the real world. The fact that the conventional approach works in simulation has misled researchers to think that it is valid in the real world.

The process of perceptual control has a technically precise, yet simple, definition (see **Perceptual Control**). What we want to perceive is compared to what we currently perceive and any difference (error) between the two drives action within the world to affect and bring the perception into line. This simple, negative feedback control process ensures that goals are achieved and maintained without requiring complex computation or explicit models of the physical world. The process is inherently adaptive and counteracts the effects of unknown disturbances. Even more significantly, when arranged in a hierarchy (**HPCT**) the same process can, the theory goes, account for all types and levels of behaviour (see **Perceptual Neural Network** and **Levels**).

DARPA

Robots entered in the DARPA challenge are clumsy and extremely slow. For from being autonomous they are partly controlled, remotely, by their human overlords.

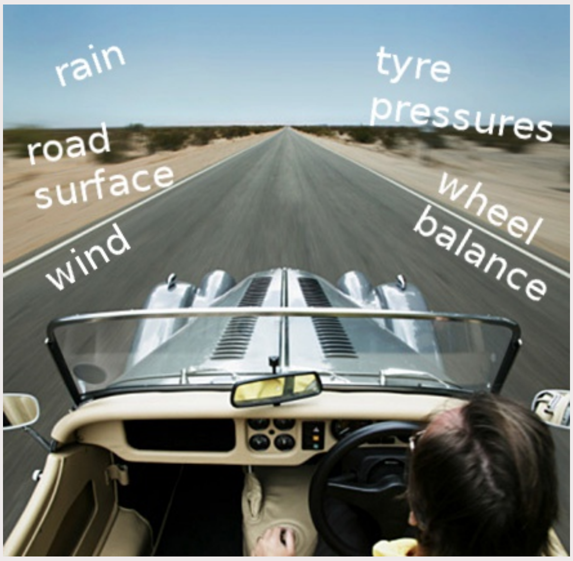

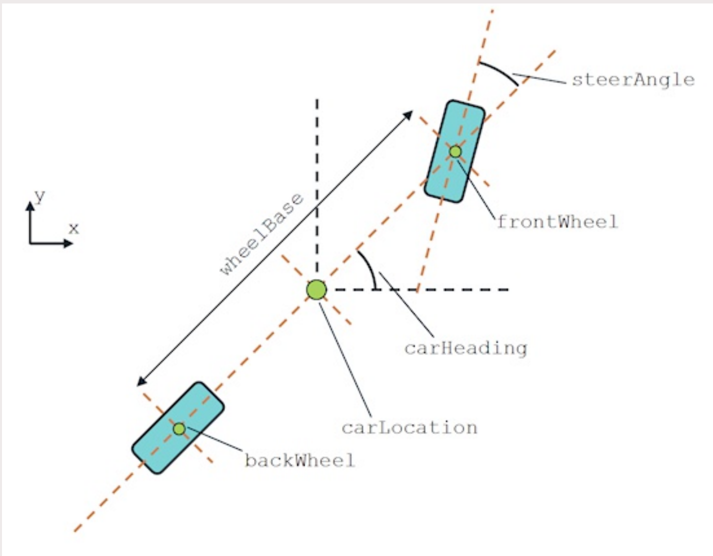

The Simulation Problem

An everyday task such as driving takes place in a dynamic, unpredictable environment. To compute the transfer function from the steering wheel to the car heading would need to take all disturbances into account which would require some sort of Laplacian Demon that knew the entire state and dynamics of the universe. Simulations of the scenario falsely give the impression that the approach is valid.

In the 19th century Pierre-Simon Laplace famously articulated the demon as an entity that could calculate all future states of the universe, but would need to know the location and momentum of every atom in the universe.

Cybernetic Control

a) The standard cybernetic model mis-applies the terminology of control theory to living systems. The incorrect interpretation is that the organism transforms input into output. b) When correctly denoted in the diagram it is seen that it is actually sensory input that feeds in to the comparator and is the variable under control. Additionally what was labelled "input" is the reference goal and is internal to living systems. The model is shown more clearly in panel **Perceptual Control**.

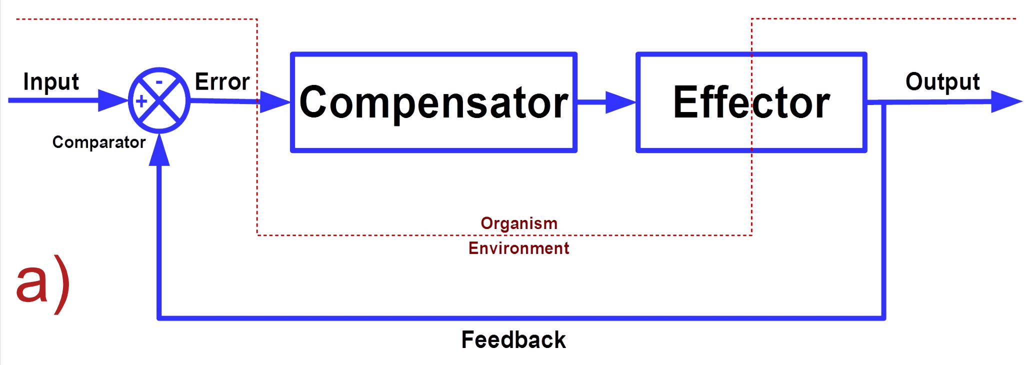

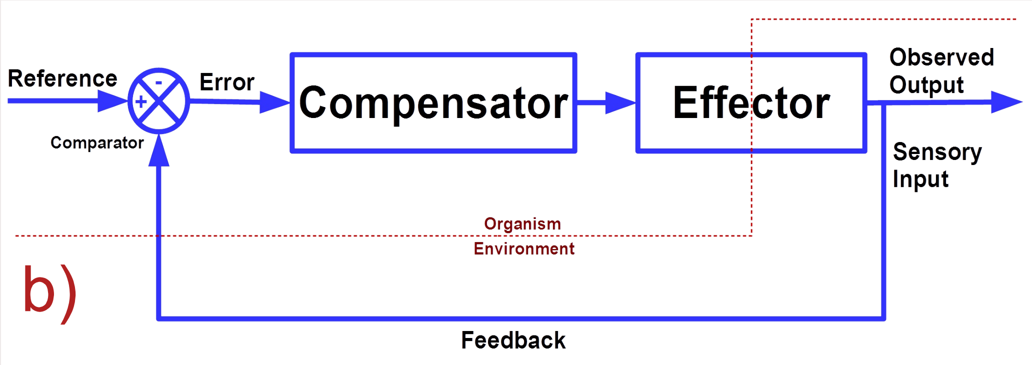

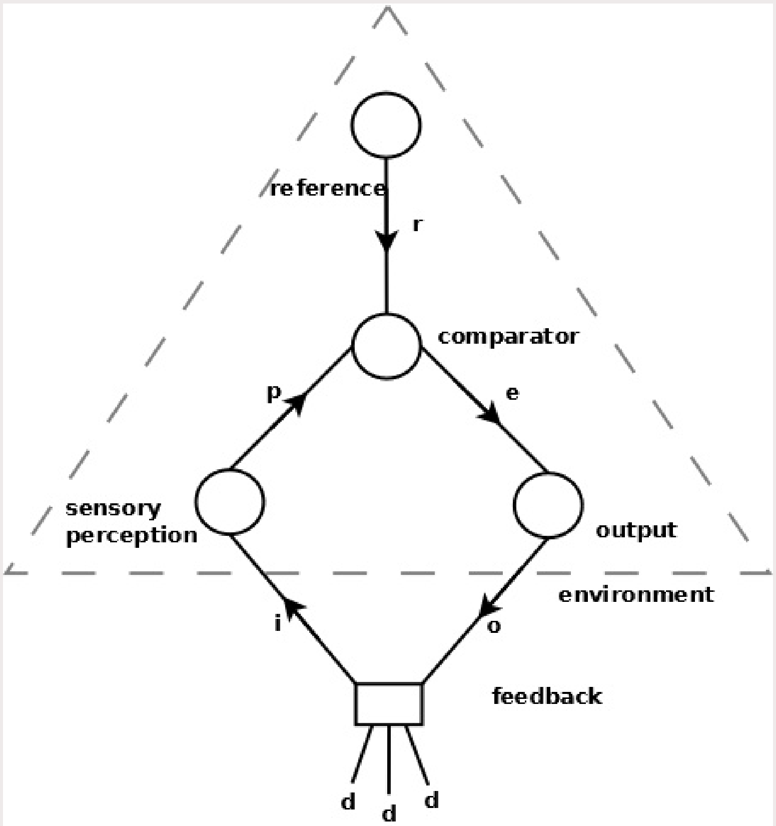

Perceptual Control

The basic perceptual control unit. A continuous negative feedback control system; *p* = current perception, *r* = goal perception, *e* = error and *o* = output; *i* is the input to the control system, which is a combination of the effects of the output of the system and any environmental disturbances, which are denoted by *d*. The area below the triangle is the feedback path, through the environment.

Perceptual Neural Network

Going up the hierarchy a perception at each level is a combination of perceptions from lower systems. Perceptions become more complex and abstract as you go up the hierarchy. Coming down the hierarchy a reference goal is a combination of output signals from higher systems, but these are not specifying commands of what to do but specifying what to perceive, and then that system will achieve its perceptual goal by varying its own outputs. Nowhere in this hierarchy are actions being selected or specified.

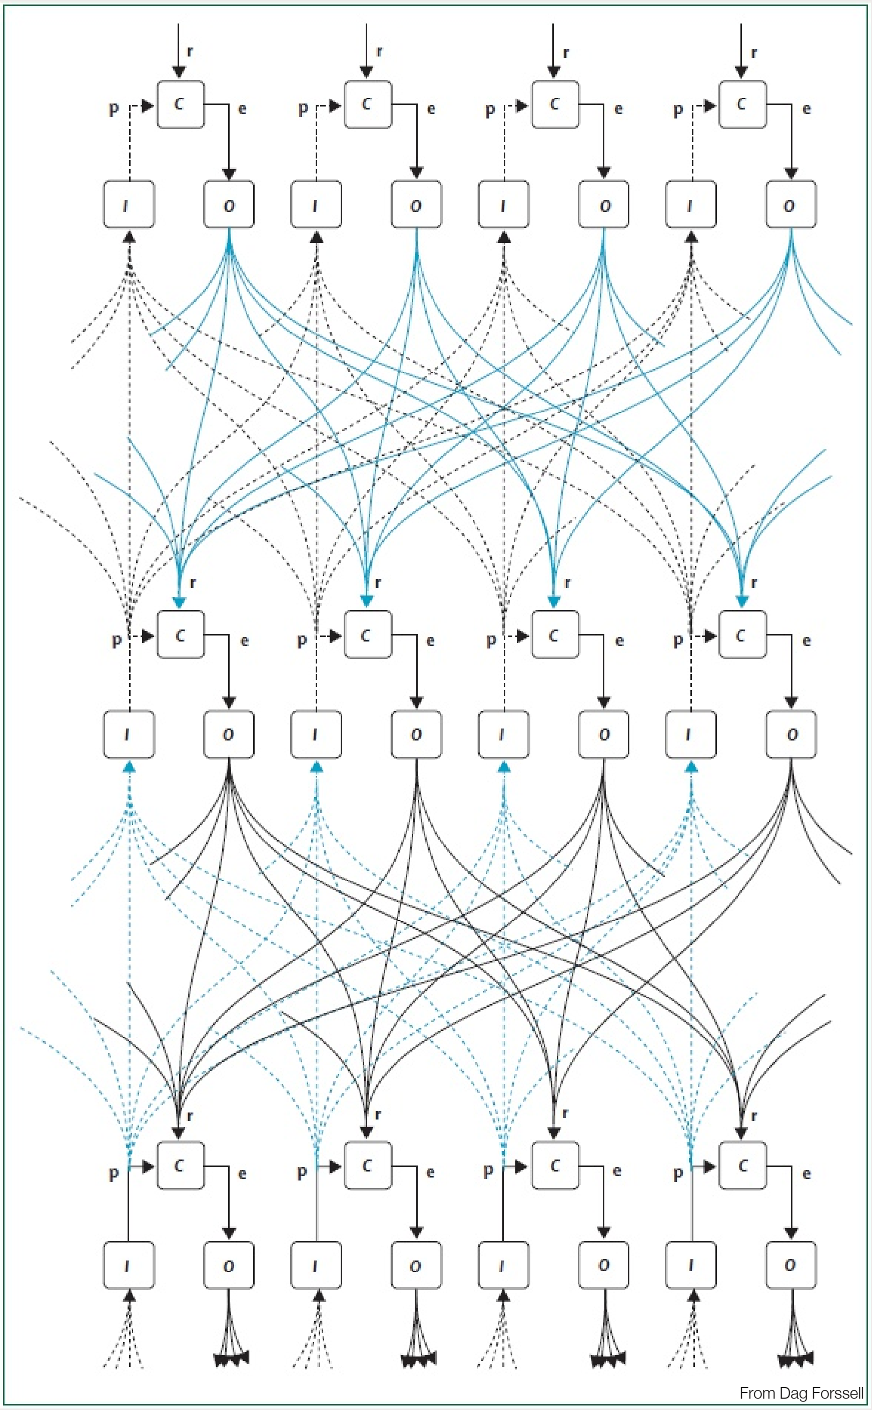

From Dag Forssell

Here are some examples that should shed light on how this simple process applies across the board as an explanation of the underlying process of behaviour:

- At the lower levels, the iris system in the eye controls the intensity of light falling on the retina, by varying the aperture.
- We control the visual relationship between our hand and an object when we want to pick something up.
- When we are speaking face-to-face with someone we control the distance between us at a comfortable level. And different people will have a different sense of what is comfortable to them.
- We can perceive and control the sequences of things, such as the letters in a word, or the toppings on a pizza.
- We can perceive events, that extend over time, such as writing a sentence, catching a train or giving a presentation.
- Cooking a meal, or performing long division, is a process of controlling a perception of a program-type process.
- At an even higher level we control more abstract perceptions. We control our own particular sense of honesty, by robbing or not robbing a bank. By voting we are acting, in a limited way, to control the concept of a desired political system.

Any complex task can be decomposed, not into behaviours, as per convention, but into a hierarchy of perceptual goals where each is represented by a simple perceptual control system.

All these are variables we can perceive, and, potentially, control. And for each there will be a value that we want to get to. Hundreds, or millions, of perceptions being controlled at any one point in time, some lasting only moments, others a lifetime.

This perceptual control architecture provides the blueprint for how to design robots which are truly autonomous and able to dynamically counter the uncertainties of the real world.

Professor Rodney Brooks, a major figure in Behavior-based Robotics and Artificial Life, has recognised that current methodologies are not sufficient for explaining the general behaviour of living systems. Perceptual control holds open the promise of furnishing Brooks' missing 'new stuff' by defining a simple, common process, a scalable, dynamical hierarchy, and organisational principle to explain how complex and intelligent behaviour emerges from simplicity.

The implications of PCT are profound not only for robotics, but also for all behavioural sciences. Research in other disciplines is slowly,

Levels of Perception

A rough idea of the types of perceptions a human controls at the different levels. Each perception is formed by a combination of perceptions from the level below.

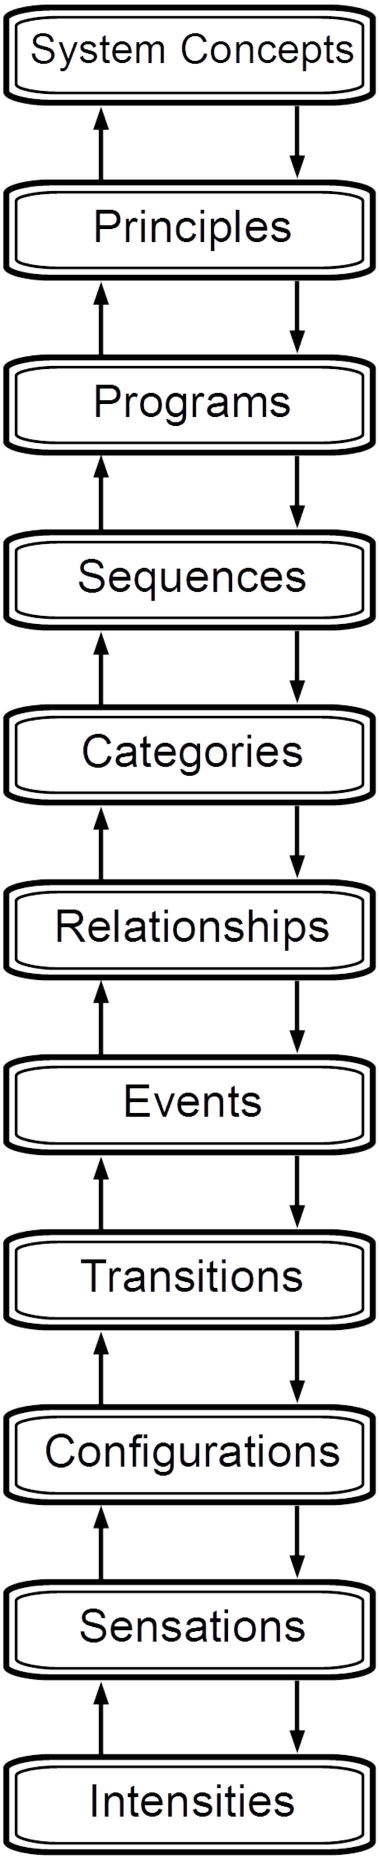

Organisational entities; e.g. religion, self, government, physics.

Life conventions; e.g. neatness, honesty, values, morality, laws.

Conditional sequence perceptions; e.g. recipe, language, mathematics.

Order of perceptions over time; e.g. letters in a word, toppings on a pizza, musical melody.

Classes of perceptions; e.g. different languages, musical styles.

Relationships between other perceptions; e.g. on, near, proximity, greater than.

Temporal arrangement of other perceptions; e.g. catching a train, bounce of ball.

Change; e.g. motion, musical scales, speed of car, rhythms.

Physical arrangement of other perceptions; e.g. musical chord, objects, faces.

Combination of sensory stimulation; e.g. tastes, temperature, pressure.

Perceived degree of sensory stimulation; e.g. brightness of light, muscle tension.

Based on Fred Nickols

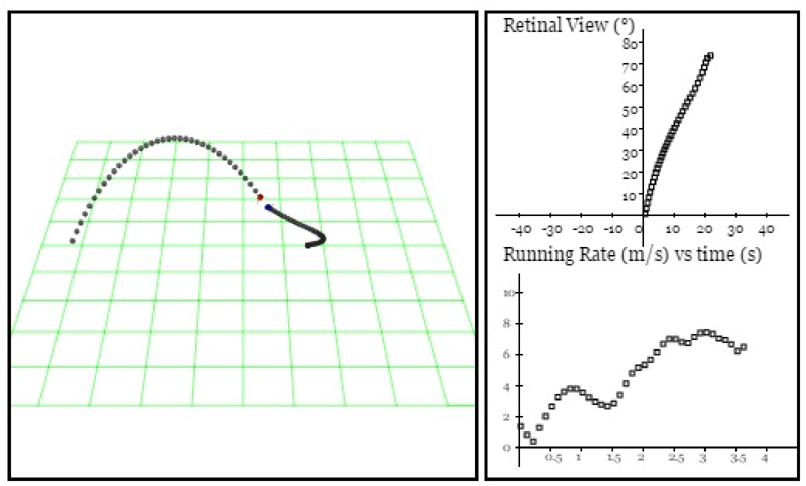

The Prediction Fallacy

The path the fielder takes to catch a baseball is determined by keeping the optical velocities of the image of the ball on the retina at constant values, rather than by predicting the trajectory of the ball and computing the route of the fielder. Try out the demo for yourself at goo.gl/jBOcmQ.

but systematically contradicting time-honoured preconceptions and methodologies.

Behaviour

An elegant demonstration of behavioural modelling from the percetual control perspective by Dr Richard Marken dismantles the assumption that predictive computation is required for seemingly complex behavioural tasks (see **The Prediction Fallacy**). Running to catch a baseball is achieved simply by controlling the position of the retinal image of the ball rather than the complex computation of trajectories and intercept points.

More generally Marken questions the basis of psychological and behavioural research arguing that the misinterpretation of the underlying processes leads to a **behavioural illusion**, from which invalid conclusions are made regarding the relationship between the environment, the organism and action.

"The behavioral illusion refers to the fact that the disturbance-resisting actions of a perceptual control system will appear to an observer to be a reaction to stimuli, particularly when the aim of these actions -- keeping a perceptual variable in a reference state -- goes unnoticed or is ignored. The implication of this illusion is that people who are trying to understand or reproduce the behavior of living systems are likely to conclude that behavior is caused or guided by outside events when, in fact, behavior is a process of controlling perceptual inputs," says Marken.

Memory

Within the hierarchy the values of the reference goals come from higher levels, which highlights the role of memory as stored perceptions, to be reused later as references. The practical result of this functionality is that control becomes more efficient.

For example, the first time you drink tea you may add sugar, bit by bit, repeatedly tasting, to control your desired perception of sweetness. It would be laborious, and impractical, if you had to repeat his process every time you drank tea, so you remember your perception of adding three spoonfuls, say. Next time you drink tea you control the desired sweetness by adding three spoonfuls of sugar without having to taste it.

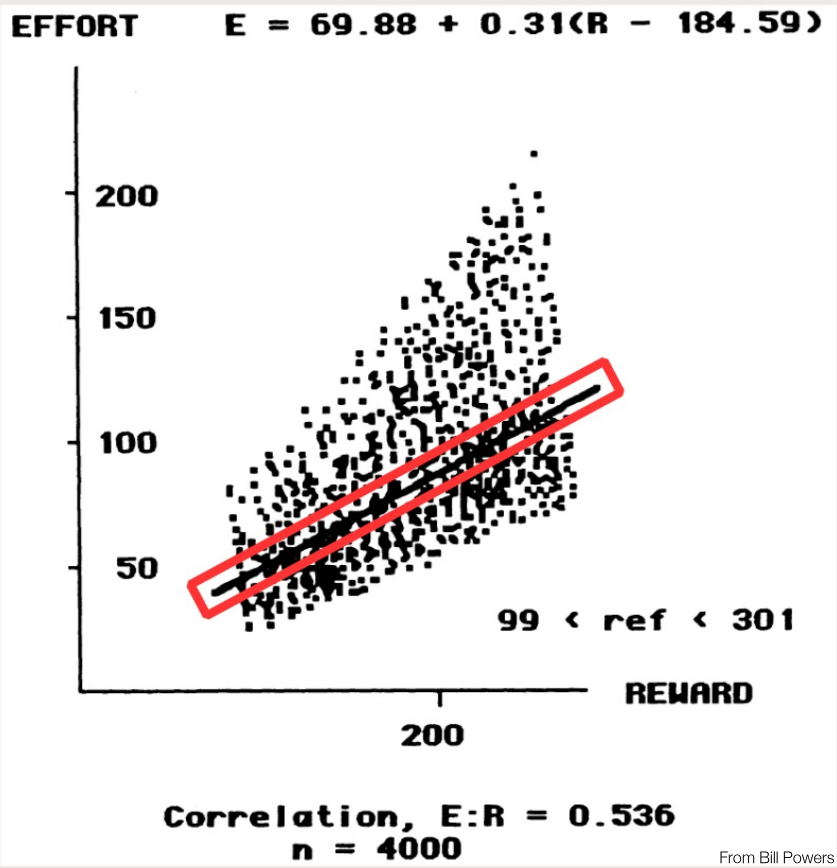

From Bill Powers

Damn Statistics

Statistics derived from the whole group incorrectly conclude a relationship between effort and reward with a slope of 0.31 and a correlation of 0.54

Learning

Control systems are not born with the ability to be able to control variables, at all levels. Like anything else the abilities need to be acquired. Initially, according to the nature of the organisation of the nervous systems, the actions you are able to take may be ineffective against the variables you are attempting to control. Furthermore, you may not have the perceptual apparatus in place to be able to perceive things which are affecting you.

PCT postulates that at a fundamental level human control systems continually perform to keep certain *intrinsic* physiological and biochemical variables at particular values. Such variables include body temperature, blood glucose levels and carbon dioxide levels. When these variables are not at their correct values then *intrinsic error* is experienced. Whenever intrinsic error persists re-organisation within the nervous system takes place. That is, connections between nerve cells are altered. This reorganisation affects the outward behaviour of the organism which, in turn, affects the perceptions and intrinsic variables. If the effect of the behaviour of the new structure of the nervous system does reduce the error, then any re-organisation is stopped or delayed. If the re-organised structure has no effect on reducing the intrinsic error then it is reorganised out of existence. In this way the system learns to control physiology by controlling perceptions.

Neuroscience

In the context of disorders neurobiological research by Henry Yin at Duke University, North Carolina has indeed found that the nervous system is organised as a hierarchy of negative feedback control systems. Adverse changes to control parameters, caused by neurological damage, can explain symptoms in movement disorders such as Parkinson's and Huntingdon's disease.

Conventional approaches to neuroscience research have not led to the expected advances. Yin says, "In the last 50 years there has been little progress in understanding how the brain works as a whole. Neuroscientists completely accepted the dogma that the organism is an input-output device that transforms sensory inputs into motor outputs, i.e. behavior. Consequently, nearly all experiments were either incorrectly designed or incorrectly interpreted, or both. Decades of experimental efforts were largely wasted. Due to repeated failures to understand the neural basis of behavior, many neuroscientists no longer believe it is possible to do so. The types of models proposed to explain behavior in neuroscience are largely fantasies, rather than working models that can be run and tested. Not surprisingly, neuroscience has largely been ignored by engineers and roboticists. By offering the first valid working model of behavior, PCT will eventually transform neuroscience and make it possible to understand how the brain generates behavior."

Mental Health

Warren Mansell at the University of Manchester applies the principles of PCT to research in clinical psychology and identifies it as a unifying theory for behavioural studies stating, "To date, the field of mental health has been flooded by a diverse and confusing mixture of different psychotherapies. PCT has common principles - control, conflict and reorganisation - that guide any effective intervention and have been shown to integrate the diverse forms of psychotherapy."

With regard to psychological problems, Mansell particularly pinpoints the cause as internal conflict between perceptual control systems. A desire to eat cake will conflict with a desire to lose weight. In general the resolution to such conflicts can only take place at a higher level, which sets those goals. In this case a higher goal of a healthy lifestyle would lower the cake-eating frequency goal thus prioritising the goal of losing weight and resolving the conflict.

A transdiagnostic cognitive therapy based upon PCT, called **Method of Levels**, has been developed, largely by Tim Carey of the Centre for Remote Health, Alice Springs, Northern Territory. The purpose is to rapidly identify and resolve conflict thus restoring control to the person in distress.

Mansell and colleagues are currently piloting an online artificial psychotherapist, **MYLO**, based upon PCT with the aim of initiating beneficial psychological change. Early findings

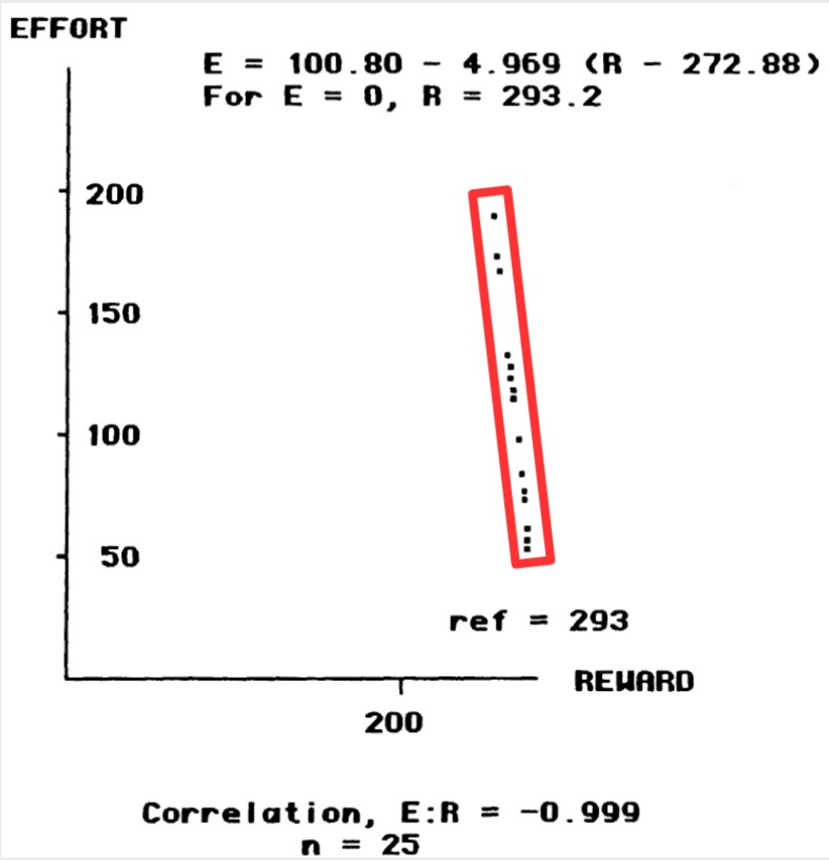

However, when the internal motivations of individuals are taken into account the true relationship emerges, showing a slope of -5,0 and a correlation of -1.0 (negative correlation).

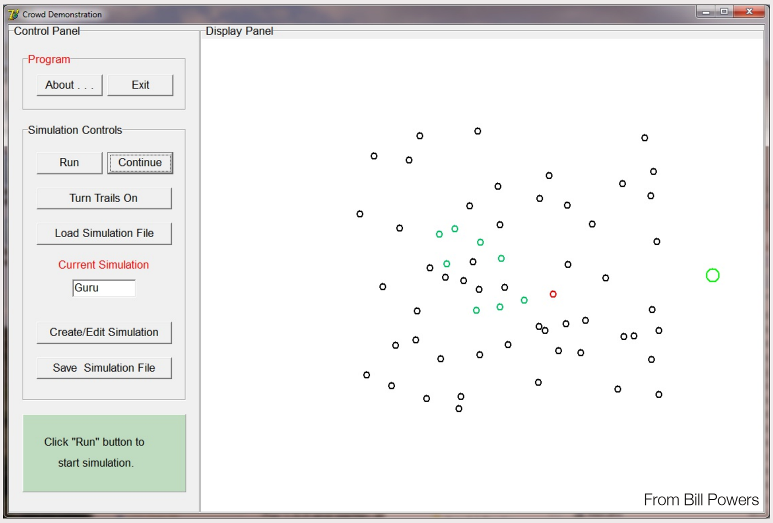

Crowd Behaviour

The behaviour of individuals within groups is governed by the simple and parsimonious process of perceptual control. Individuals easily find their way through a complex environment by controlling the perception of their proximity to each other, objects and to a target. (See video youtu.be/SR_hZw2_5YI).

indicate that its use does lead to reductions in depression, anxiety and stress.

Ethology

In the arena of animal behaviour studies, Heather Bell at the University of California, San Diego, has shown that rats control the perception of space between themselves and others when defending food. The same process was found in other species, separated by millions of years of evolutionary history, suggesting that it is likely to be a common, general process throughout the animal kingdom.

Sociology

When we realise that individuals are perceptual control systems it can be seen that there is potential for the theory to impact Sociology, to provide explanations of how individuals interact within groups or crowds (see **Crowd Behaviour** demo). Kent McClelland at Grinnel College, Iowa has theorised *Collective Control* as a new and better conceptual foundation for understanding social structure and culture. Additional work in the field finds that people are controlling their own perceptions when using language and interpreting social events, and that a person's social identity can be understood as a controlled perceptual variable.

Statistics

The conventional wisdom of behavioural studies is that the larger the sample size the better. That idea needs to be ditched, if the different internal motivations of the individuals are not taken into account. Simulated data (see **Damn Statistics**) plots effort against reward for 4000 individuals. The statistics for the whole group indicate that greater effort reaps greater reward. At first glance, that appears to make perfect sense. However, when we look at individuals, or rather sets of individuals with the same goal, we see a very different picture. That is, to achieve the *same* goal different individuals within the set need to exert *different* effort, depending upon their circumstances. So, for example, if a set of individuals all want to buy a BMW, for $50,000 say, those who have lower salaries or more outgoings will have to make a greater effort than those on a higher salary. To appreciate the actual relationship between behavioural variables it is necessary to analyse people as individual perceptual control systems.

Planning

Normally planning is thought of as a fairly complex process of computing and reasoning a solution within a sophisticated model, based upon extensive pre-defined knowledge.

Robotics in the Real World

Robot implementations based upon the principles and architecture of perceptual control.

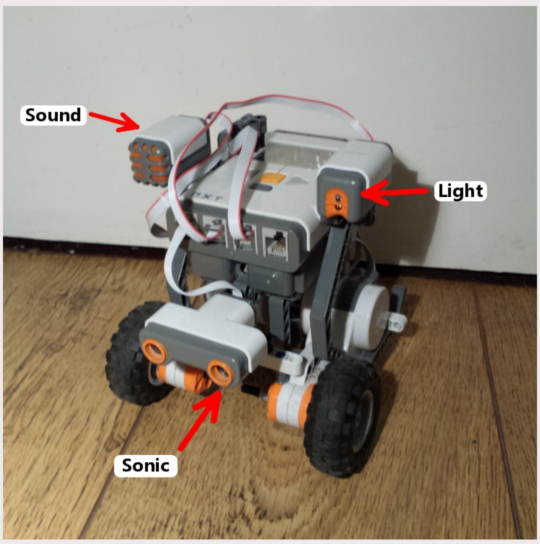

Complex behaviour emerges from a basic robot system embodying a perceptual control hierarchy (youtu.be/xtYu53dKz2Q).

Robot performs Tai Chi movements on-the-fly by simultaneously controlling perceptions at multiple levels rather than by pre-defined kinematic poses (youtu.be/05OpquGG7AQ).

With PCT insight, though, it can be seen in relatively simple terms as the same process of perceptual control, except that it is done in imagination. So, planning is a matter of thinking through what you would need to perceive to perform the task, but without actually carrying out any action. The role of knowledge is fulfilled by perceptions as memory. For example, if I want to plan a route to the motorway from my house, I think of the perceptions (not the actions) that I would need to control in real life, such as perceiving myself turning left at the traffic lights, turning right at the roundabout then bearing left onto the slip road. Although these could be thought of as perceptions of actions the crucial point is that it is the perception of them that informs us that we are successfully executing them. It is the perception being controlled.

Reinforcement Learning

The concept of reinforcement learning needs to be re-evaluated in the context of PCT. The traditional implication is that rewards and punishment cause attractant or avoidance behaviour. However, from the PCT perspective reward and punishment only make sense if they are things that the organism wants to acquire (or avoid). The rats in Behaviourist experiments were starved beforehand, giving them the opportunity to reduce their perceptual error. If they weren't hungry, rewards would have had no effect on their behaviour. In other words, reinforcement is a misnomer and is actually the same old process of perceptual control. Behavioural scenarios are better understood by identifying the perceptual variables that are being controlled. If a child is bullied after school they may misbehave in order to be sent to detention in order to feel safe and avoid abusers. In this case the supposed effect of the detention punishment is non-existent, as an unrelated perception is being controlled.

There appears to have been a degree of success applying the reinforcement learning methodology, along with deep neural networks, to learning within AI systems, such as the playing of Atari console games. However, the approach relates input game states to output actions. In the real world different output actions are required for the same input states according to the environmental disturbances present. It remains to be seen, therefore, if the reinforcement learning methodology can extend beyond the simulated world.

Robotics

PCT provides a technical language and physical architecture for understanding the general behaviour of living systems. Potentially this provides an ideal solution for the architecture and methodology for developing artificial systems, especially given the simplicity and universality of the central process. In my Artificial Life paper I describe how the perceptual control architecture can be applied to robotics systems. For the specific robot system implemented for the paper the hardware is pretty basic; two motors and the equivalent of the two pixels for sensors. But because it embodies a relatively sophisticated perceptual control hierarchy the system exhibits complex behaviour, where it controls, for example, a perception of a sequence of events, searching for maximum light, avoiding obstacles and correcting its direction. It also controls its perception of the change within internal variables, by which it gets itself out of deadlock situations.

The perceptual control architecture is a general organisation that can be applied to any area of robotics as well as, in principle, providing a blueprint for the attainment of the holy grail of AI research, psychologically-advanced and intelligent systems.

Self-driving cars are a hot topic at the moment and are seen as an imminent manifestation of artificial intelligence. However, there are good reasons for being wary of the hype. Essentially the methodology comprises the geometric manipulation of an object (the car) within a 3-dimensional model-space comprising other vehicles, buildings and people. The creation of the model itself requires significant external support by way of detailed *manual* mapping of the driving environment. But there is more to driving than low-level object manipulation. It also requires dynamic response to novel circumstances as well as psychological interpretations of the intentions of other control systems (that's other people to you and me). So far, self-driving cars have largely been restricted to test environments. Unless the systems are able to incorporate understanding and control of high-level, psychological perceptions then that is where they will remain.

Consciousness

Can PCT tell us anything about why we are conscious? Well, perhaps the

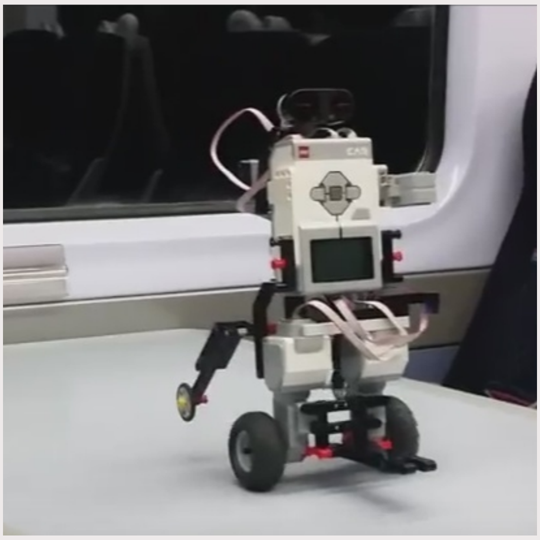

A robot on a speeding bumpy train stands up and balances by itself compensating for unknown disturbances (youtu.be/FCPDEeosCPU).

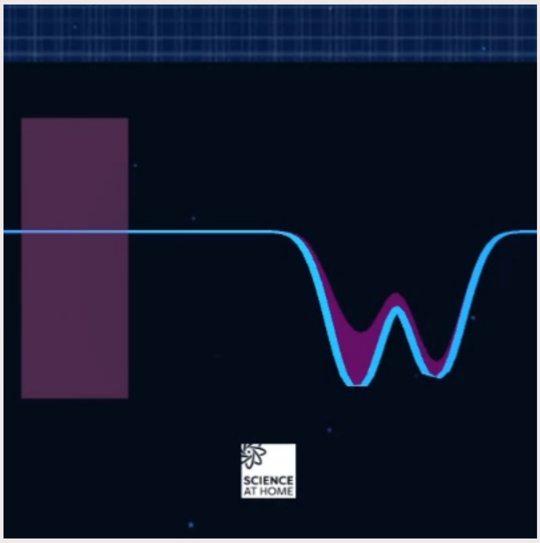

Perceptual control agents play the Quantum Moves game manipulating an unstable atom (youtu.be/0h3KPqWpRuA).

more interesting and relevant question is 'why are we *not* conscious?'

Most of what we do does not impinge on conscious awareness. There are some things of which we are *never* aware; the light-controlling iris system in the eye, for example. Why are we never conscious of this system operating?

There are some things of which we were once aware but are no longer. When learning a new task or skill, such as playing a guitar or walking, we focus our attention on the minutiae of what is required to pick up the skill, but once, after practice, we have learned the skill we are no longer aware of the details. When learning to play a basic guitar note we may, initially, concentrate on, and be conscious of, the pressure of a finger we are applying to a string to get a pure sound, void of buzzing. An accomplished guitarist is not conscious of such details; until a bum note is played, that is. Presumably the same was the case when learning to walk, in that we would consciously move and place a leg in such a way to achieve a particular perceived goal. Now, as a grown up, it is all second nature and we are able to walk without being conscious of the details. Why do we cease to be aware once something has been learned?

Current conscious awareness can shift in an instant, seemingly determined by external events, such as a loud noise, or self-directed by an internal stream of thoughts. What was not conscious can come into awareness, replacing a previous focus. And what was conscious is lost.

There is a condition called Blindsight where people have no conscious awareness of being able to see, however when asked are able to identify the object in their field of view. Also, they may catch an object thrown to them though they claim not to be able to see it.

These highlight that there is an association between the presence of error within control systems and awareness. When a system becomes finely-tuned and error free then conscious awareness disappears. Arising from this perspective the role of consciousness is to aid learning by directing resources to particular systems which are in need of reorganisation for the purpose of improving the quality of control.

So, maybe we can think of consciousness itself as a different type of perception, a perception of the quality of control, and consciousness only occurs when that quality is not optimal.

Robotics in the Real World

The application of robotics has largely focussed on mechanical manipulation of objects in geometric space; navigation, pick-and-place, grasping, autonomous vehicles, assembly robots. Artificial Intelligence has focussed on high-level reasoning and planning. Each discipline employs a variety of techniques with little consensus or overlap between different areas. There is no underlying concept that provides a common explanation or rationale of behaviour which could act as a unifying thread to bring together and combine what should be complimentary and integrated fields.

The focus solely on mechanical manipulation may seem sensible but is actually an unconscious, self-imposed restriction on robotics due to the inadequacy of the conventional methodologies of providing an explanatory perspective of general behaviour. There are many other purposeful things that living systems do which are just ignored because conventional approaches provide no coherent understanding of the underlying processes, for example, scratching your nose, laughing, having cosmetic surgery, getting drunk, sex, gambling, tidying room, supporting gender equality, and many, many more.

"Nothing in behaviour makes sense except in the light of perceptual control."

Conventional approaches to robotics have little to say about these sorts of behaviours. However, although the objective may not be to reproduce them all in robots, the fact that they, and associated underlying processes, are neglected means that similar behaviours, which may be very useful to robots, will never materialise, severely impeding future progress in the field.

Natural Selection is the central process of evolution. Perceptual Control does provide a definition of the underlying process of purposive behaviour that applies to these behaviours and to mechanical manipulation and to high-level cognitive abilities. To paraphrase Dobzhansky, "Nothing in behaviour makes sense except in the light of perceptual control." In other words, Perceptual Control is the 'Natural Selection' of behaviour, and robotics.

Perceptual Control is being successfully applied to real world robotics systems, see panel **Robotics in the Real World**. Also shown there is the successful control, in a simulated environment, of an unstable quantum state by the PCT methodology, without the need for the modelling of that the quantum system.

For 70 years Robotics and AI have been doing things the hard way. The result has been that beyond simulated or controlled environments the conventional, computational, approach becomes mired in complexity. PCT takes a lesson from nature and shows that, for complex problems, even in complex environments, a hierarchy of perceptual control systems resolves and dissolves the computational problem.

Some of the benefits of perceptual control systems are:

- **Simplicity** - The basic unit of perceptual control systems has a very simple operation.
- **Universality** - The basic perceptual control process is common to all levels and types of behavior.
- **Scalability** - The arrangement of basic control units into an interdependent hierarchy results in a highly scalable architecture.
- **Adaptivity** - The structure of the basic perceptual control unit is inherently adaptive. A disturbance to the perceptual input results in error, which in turn results in output that acts upon the perception, automatically canceling out the effects of the disturbance.
- **Autonomy** - As the goals of the system and the means by which they can be achieved are themselves embodied within a perceptual control system, it can be said to be truly autonomous.
- **Complexity** - The power of HPCT and the complexity of behavior derive not from complicated objective models of the physical world, but from the control of high-level, sophisticated subjective perceptions, along with the multitude of lower-level perceptions on which they depend.

The implications for robotics of these benefits are profound, resulting in robots that are likely to be simpler, cheaper and significantly more sophisticated than existing systems.

Initial implementation of Perceptual Control Theory is at a low-level, but the architecture is inherently scalable and the theory provides a roadmap of how to work from the bottom-up to build ever more sophisticated systems. With the understanding of the approach in the context of psychology and general behaviour it looks more promising that, at last, robotics can move out of the controlled conditions of the factory floor into the unpredictable, dynamic real world to emulate the intelligent, autonomous and psychologically-advanced behaviour of humans and other living systems.

Websites

Perceptual Robots

www.perceptualrobots.com

PCT Web

www.pctweb.org

International Association of Perceptual Control Theory

www.iapct.org