

PCT and Beyond: Towards a Computational Framework for 'Intelligent' Communicative Systems

Roger K. Moore

Department of Computer Science, University of Sheffield
Regent Court, 211 Portobello, Sheffield, S1 4DP, UK

r.k.moore@sheffield.ac.uk
<http://www.dcs.shef.ac.uk/~roger>

Abstract. Recent years have witnessed increasing interest in the potential benefits of 'intelligent' autonomous machines such as robots. Honda's *Asimo* humanoid robot, iRobot's *Roomba* robot vacuum cleaner and Google's driverless cars have fired the imagination of the general public, and social media buzz with speculation about a utopian world of helpful robot assistants or the coming robot apocalypse! However, there is a long way to go before autonomous systems reach the level of capabilities required for even the simplest of tasks involving human-robot interaction - especially if it involves communicative behaviour such as speech and language. Of course the field of Artificial Intelligence (AI) has made great strides in these areas, and has moved on from abstract high-level rule-based paradigms to embodied architectures whose operations are grounded in real physical environments. What is still missing, however, is an overarching theory of intelligent *communicative* behaviour that informs system-level design decisions in order to provide a more coherent approach to system integration. This chapter introduces the beginnings of such a framework inspired by the principles of Perceptual Control Theory (PCT). In particular, it is observed that PCT has hitherto tended to view perceptual processes as a relatively straightforward series of transformations from sensation to perception, and has overlooked the potential of powerful generative model-based solutions that have emerged in practical fields such as visual or auditory scene analysis. Starting from first principles, a sequence of arguments is presented which not only shows how these ideas might be integrated into PCT, but which also extend PCT towards a remarkably symmetric architecture for a needs-driven communicative agent. It is concluded that, if behaviour is the *control* of perception (the central tenet of PCT), then perception (at least for communicative agents) is the *simulation* of behaviour.

Keywords: perceptual control theory, artificial intelligence, cognitive systems, robotics, communicative agents

1 Introduction

The past few decades have seen an enormous growth in the level of interest being shown in so-called ‘intelligent’ machines [1]. Funding agencies and large corporations worldwide have been investing heavily in autonomous systems - especially *robotics* - on the premise that significant economic benefits may be derived from automating hitherto people-intensive activities. Ranging from production-line robots to homecare assistants, and from robotic surgeons to driverless cars, future intelligent systems are expected to transform our lives in much the same way that the invention of the steam engine accelerated the industrial revolution during the 18th and 19th centuries [2,3,4,5,6].

However, in order to deliver the expected benefits, future autonomous systems will indeed need to be intelligent - they must integrate seamlessly into real-world environments, act appropriately in complex physical and temporal situations, solve difficult logistical problems, interact effectively with human users/operators using accepted social conventions (such as speech and language), be robust in the face of unpredictable disturbances and interruptions, operate independently within an accepted ethical framework, and be at least partially responsible for their own behaviours [7]. This is a challenging wish list that goes well beyond the bounds of traditional fields such as Artificial Intelligence (AI). Indeed, the requirements for intelligent systems/robots are so demanding that insights need to be integrated from a wide array of disciplines ranging from engineering and computer science to psychology, cognitive neuroscience and linguistics.

In practice, a robot is a complex physical entity often consisting of a large number of moving parts, an electrical power system, an array of electronic components including on-board and off-board processors, wired and wireless communication links, various computer operating systems and a range of software modules for managing the overall system. This means that it is not only necessary to develop comprehensive tools and techniques for designing, building and programming such devices to meet particular application requirements, but it is very likely that approaches will also need to be based on a deeper understanding of how existing intelligent systems - living organisms - solve the challenges listed above [8,9].

At present, the sheer complexity of such systems, coupled with the high cost of developing bespoke hardware platforms, inevitably means that components are integrated without a great deal of thought being given to more general principles of intelligent behaviour. Off-the-shelf solutions for locomotion, navigation, manipulation and interaction are often combined with independent modules for input/output modalities such as vision, speech and gesture without too much consideration of potential synergies. As a result, behaviours may be programmed in an ad-hoc manner using heuristic (rather than principled) approaches to the

necessary algorithms¹. What is required is an overarching theory of intelligent behaviour that informs system-level design decisions in order to provide a more coherent approach to system integration, especially for communicative agents.

1.1 Good Old-Fashioned Artificial Intelligence

The term Artificial Intelligence (AI) was coined in 1955 and, in its early years, was mainly concerned with mathematical logic and automatic theorem proving (on the assumption that ‘intelligence’ was founded on processes of rational thought) [10]. Activities such as playing chess were seen as the epitome of human intellectual achievement, and thus became the focus of early research. However, it was soon realised that physically manipulating actual chess pieces could be more challenging than playing the game itself. AI thus moved on to fulfil an important role in expanding our understanding of how living systems in general, and human beings in particular, interact physically with the world through developments in areas such as bio-inspired robotics and autonomous systems.

Until the mid-1980s, the main paradigm for AI-based robotics was the so-called ‘deliberative’ architecture in which symbolic representations of the world were manipulated in a hierarchical rule-based framework involving goals and subgoals. The process operated using a *Sense* → *Plan* → *Act* cycle (see Fig. 1), very much in tune with the ‘behaviourist’ *Stimulus* → *Response* (S-R) framework that was dominating the field of psychology at the time [11].




Fig. 1. Illustration of the sequential deliberative (*Sense* → *Plan* → *Act*) architecture popular in Good Old-Fashioned Artificial Intelligence (GOFAI).

When applied to robotics, the deliberative AI paradigm essentially took a high-level approach in which functions such as navigation were regarded as problem-solving activities that required the manipulation and search of appropriate symbolic data structures. However, it soon became apparent that this perspective - subsequently termed Good Old-Fashioned Artificial Intelligence (GOFAI) [12] - suffered from severe limitations, particularly with respect to an over-reliance on accurate world models and an inability to respond quickly to changing situations and context.

¹ Of course it’s not that such theory doesn’t exist. Rather, the demands of contemporary intelligent systems are such that it is often necessary to take a pragmatic approach to system implementation.

1.2 Behaviour-Based Robotics

By the 1980s, the mounting difficulties faced by attempting to deploy GOFAI in real-world robots had led to a new ‘reactive’ (as opposed to ‘deliberative’) approach based on low-level *representation-free* processes. Hailing the establishment of what was to be known as ‘New AI’, Rodney Brooks introduced a novel ‘subsumption’ architecture in which the emphasis was on real-time behaviour using simple computations embedded in a layered structure [13], [14] (see Fig. 2). The basic idea was that the different layers should operate more or less independently (and in parallel), with the higher levels relying on successful operation of the lower levels. If necessary, the higher levels could change (that is, ‘subsume’) the behaviour of the lower levels. Overall, the emphasis was on the *grounding* of behaviour in the real-world, with the hypothesis that complex interactions should arise as an *emergent* property of an array of simple processes (as had been elegantly proposed by Braitenberg in the 1980s [15]).

The subsumption approach initiated a trend towards what became known as ‘behaviour-based robotics’ [16], and the emphasis shifted from high-level abstract problem solving to low-level grounded intelligence.

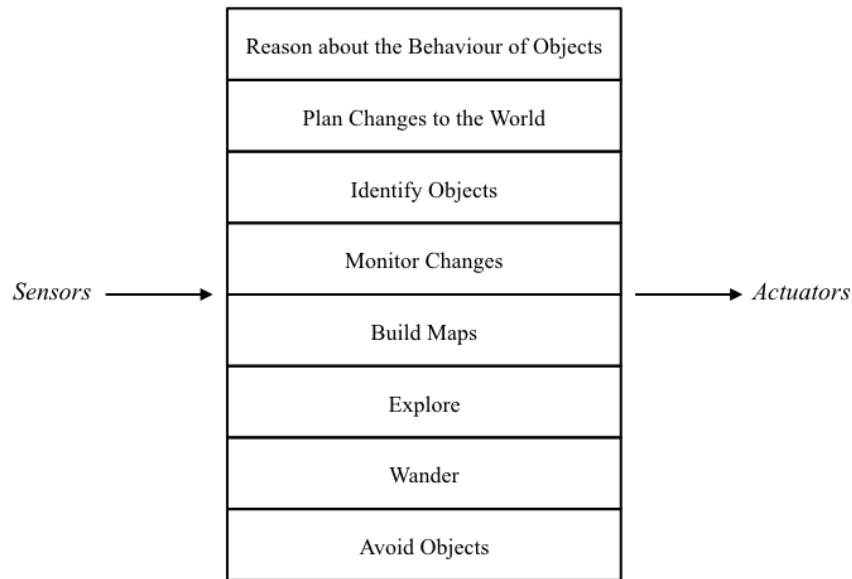


Fig. 2. Illustration of the layered reactive subsumption architecture popular in behaviour-based robotics.

Behaviour-based robotics was sufficiently successful in solving practical real-world problems that, in 1990, Rodney Brooks founded the company iRobot[®] to develop and market the Roomba[®] vacuum cleaning robot. More than 10 million Roomba robots have been sold worldwide, making it the most successful mass-market robot to date. However, the subsumption approach has been criticised for being difficult to scale up to more complex robots. Also, the emphasis on representation-free processing, whilst being popular with a subset of psychologists, is ultimately a significant restriction².

1.3 Artificial Cognitive Systems

Recent work in robotics has taken a more pragmatic view by mixing and matching the best ideas from GOFAI and behaviour-based robotics with some of the new perspectives emerging from the field of cognitive neuroscience. Known as ‘artificial cognitive systems’ [17,18], the most significant influence has been the emphasis on enaction, embodiment and situatedness in which the relationship between a goal-driven robot and its physical real-world context is paramount [19,20]. Low-level interactions are managed locally and high-level representations are given ‘meaning’, not by hand-crafted rules, but by virtue of their grounding in real-world interactions. Interaction is facilitated by the ‘affordances’ [21,22] provided by a robot’s environment.

A particular example of a contemporary architecture for an artificial cognitive system is Distributed Adaptive Control (DAC) [23,24]. The DAC architecture is based on the hypothesis that the living brain maintains a stable relationship with its environment by continuously solving the *How, Why, What, Where, When* (H4W) problem. Implementation uses an artificial neural structure organised into soma, reactive, adaptive and contextual layers, and columns which represent exosensing (defined as the sensation and perception of the world), endosensing (detecting and signalling states derived from the physical self) and action (the interface to the world) - see Fig. 3.

Unlike their GOFAI predecessors, modern cognitive approaches to intelligent systems (such as DAC) view action and perception as being synergistic (rather than separate) processes [25], and that skills should be acquired through the robot’s own active exploration of the world (rather than being pre-programmed by the system designer) [26,27]. Such approaches very much reflect contemporary models of living systems, particularly the discovery in the 1990s of so-called ‘mirror neurons’ - neural structures which appear to provide a vital link between sensory and motor behaviour and thereby a mechanism for action understanding, imitation and learning [28,29,30].

A recent book by Murray Shanahan provides an excellent overview of the contemporary perspective in artificial cognitive systems [31].

² In reality, the proposal that internal models were unnecessary was a stance adopted by Brooks to make a strong point about the inadequacies of GOFAI. In practice, the subsumption architecture incorporates such models (for example, the building of maps is mentioned in Fig. 2).

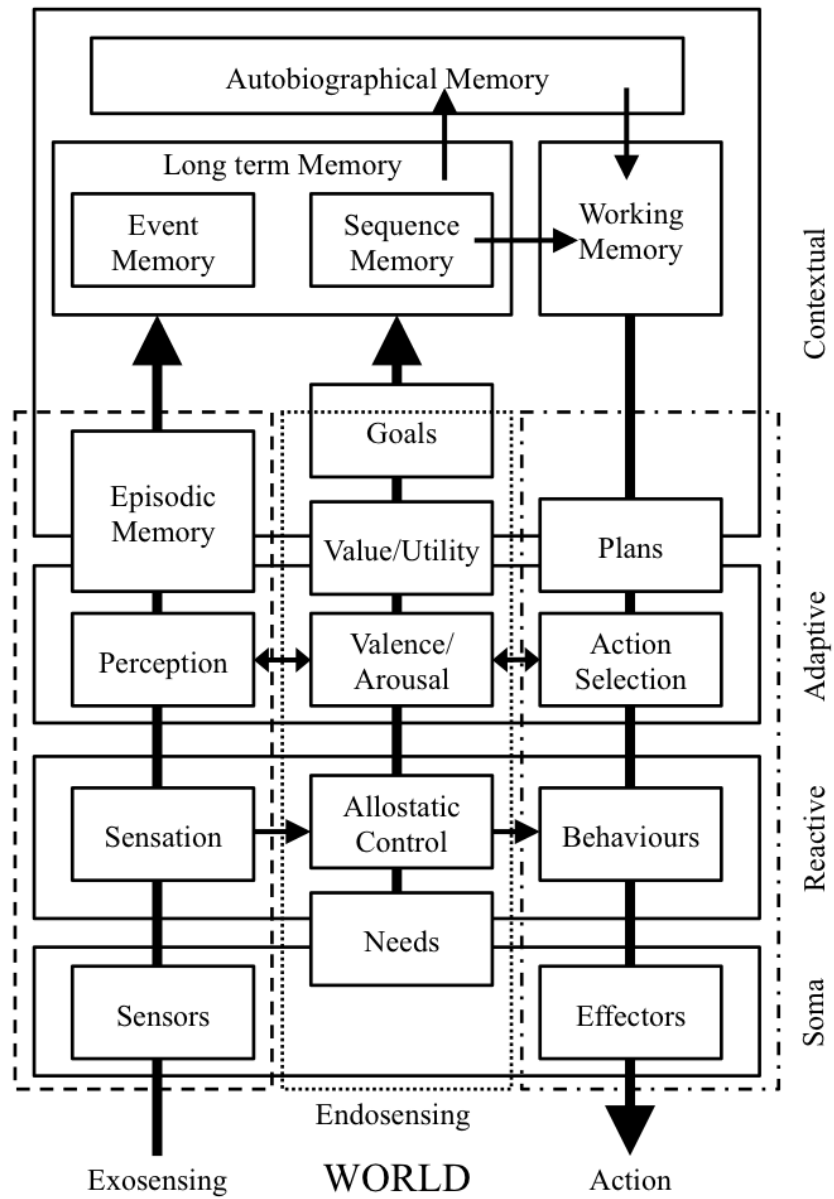


Fig. 3. Illustration of the Distributed Adaptive Control (DAC) architecture organised into exosensing, endosensing and action columns, and soma, reactive, adaptive and contextual layers (derived from Fig. 1 in [24]).

1.4 Agent Based Modelling

One aspect of behaviour that is common across alternative approaches to modelling and building intelligent systems is *intentionality*. It is clearly the case that living systems appear to be goal-directed and purposeful in their endeavours [32], and this has had a major influence on AI and robotics. In particular, one area in which intentionality plays a key role is the field of Agent-Based Modelling (ABM) [33,34]. ABM is a well-established methodology for simulating the actions and interactions of multiple agents: for example, predicting the behaviour of crowds, optimising a supply chain or managing a workforce. Various modelling paradigms are employed, such as cellular automata [35] or dedicated multi-agent programming environments [36]. However, for ‘intelligent’ agents, ABM simulations are often constructed using a Beliefs Desires Intentions (BDI) architecture [37,38] on the premise that such internal structures are required to adequately condition the behaviour of individual agents - see Fig. 4.

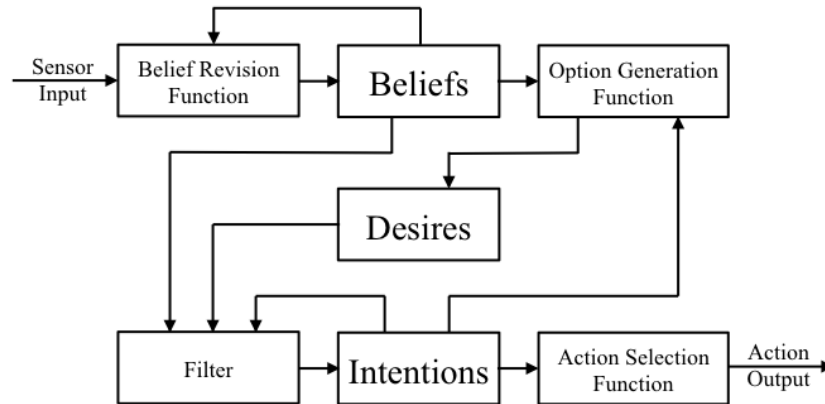


Fig. 4. Illustration of the Beliefs, Desires and Intentions (BDI) architecture used in the field of Agent-Based Modelling. ‘Beliefs’ capture the informational state of an agent, ‘Desires’ represent its motivational state and ‘Intentions’ represent its deliberative state.

BDI is a powerful approach to modelling agents, and it has been applied successfully to robotics [39,40]. However, BDI does not specify how to recog-

nise/interpret behaviour under conditions of ambiguity or uncertainty - a crucial feature of intelligent systems.

1.5 Contemporary Intelligent Systems

Over the past ten years, the field of robotics and autonomous systems has grown in stature and achievement. Robots such as Boston Dynamics' *Big Dog* and Honda's *Asimo* have successfully demonstrated that modern-day computing and electronics are finally fast enough to permit real-time control of complex behaviours such as running or climbing on uneven surfaces. Robotic quadcopters have been shown flying through moving hoops [41] and playing tennis [42], robot hands have been programmed to catch thrown objects [43] and NASA's *Robonaut R2* has been assisting on the International Space Station since 2011 [44]. Also, recent years have seen tremendous growth in the field of Artificial Neural Networks (ANNs), particularly with the success of *deep learning* as a mechanism to optimise the parameters of such multilayered networks on massive amounts of training data [45,46].

Nevertheless, despite this tremendous progress, there is still a long way to go before artificial intelligent systems will be able to demonstrate the flexibility, robustness and autonomy exhibited by even the simplest living organism [47,48]. Many contemporary robots rely on teleoperation by human operators to overcome weaknesses in their overall design, and attempts to create humanoid robots are fraught with difficulties ranging from the risk of repulsing human users due to the uncanny valley effect [49,50], to the more general uncertainties associated with human-robot interaction [51,52,53] - especially if it involves communicative agents and spoken language [54].

For example, probably the most well known communicative agent is *Siri*, Apple's voice-based personal assistant for the iPhone. Released in 2011, *Siri* suddenly brought spoken language technology to the attention of the non-specialist user. Voice dictation software for document creation on PCs had been available since the 1990s, culminating with the release of Dragon System's *Naturally Speaking* and IBM's *ViaVoice* products in 1997. The big difference was that *Siri* combined automatic speech recognition and speech synthesis with natural language processing and dialogue management in order to facilitate a more *conversational* interaction between users and smart devices, and competitors such as *Google Now* and Microsoft's *Cortana* soon followed.

In reality, the practical value of contemporary communicative agents is somewhat in doubt (as evidenced by the preponderance of videos on *YouTube* which depict humorous rather than practical interactions). This has been confirmed by a recent survey that discovered only 13% of respondents use their voice-based personal agent daily, whereas 46% had tried it once and then abandoned it [55]. One reason for the lack of usability is that contemporary communicative agents are founded on a classic *stimulus*→*response* architecture (as illustrated in Fig. 5); a user speaks, their utterance is processed, a response is formulated and the system speaks back. Such an approach completely overlooks the reality of *languageing* as an emergent property of the dynamic coupling between *intentional*

agents that serves to facilitate distributed sense-making through cooperative behaviours [57,58,59,60,61]. Furthermore, the contemporary view is that language is based on the co-evolution of two key traits: *ostensive-inferential* communication and *recursive mind-reading* [62].

These are very sophisticated notions but, as yet, there is no clear practical framework for implementing such concepts. As a result, the field is still open for new ideas, and some of them may not be so new. For example, what might Perceptual Control Theory (PCT) - established over 50 years ago [63,64] (and the topic of this volume) - have to contribute to future ‘intelligent’ communicative systems?

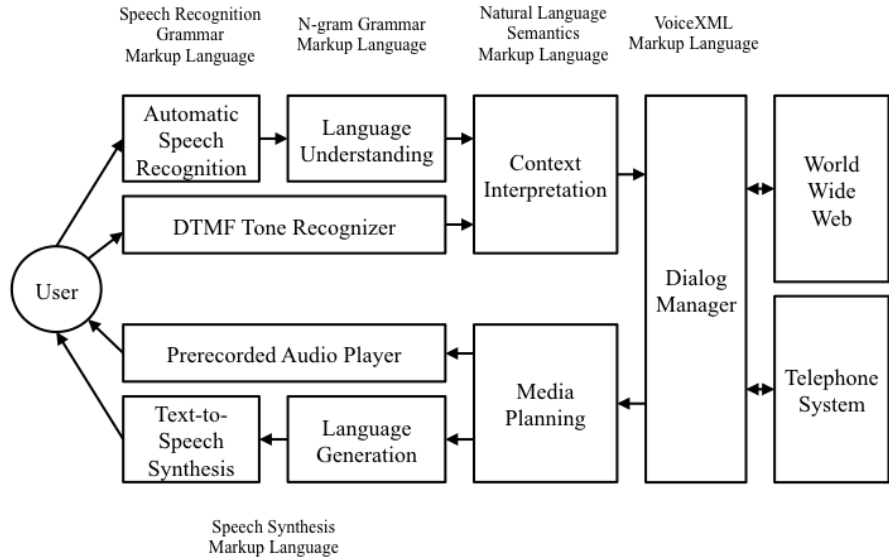


Fig. 5. Illustration of the standard *Speech Interface Framework* published by the World Wide Web Consortium (W3C) [56].

2 Whither Perceptual Control Theory?

Although it is more than 40 years since the publication of Bill Powers’ seminal book ‘Behavior: The Control of Perception’ (B:CP) [65], Perceptual Control Theory may still have something important to say about computational models of ‘intelligent’ communicative behaviour. Over the years, the main thrust

of PCT has been to question the traditional behaviourist *stimulus*→*response* stance that still prevails in some sections of the psychology field [66,67]. PCT aims to provide a more parsimonious model which highlights the critical importance of acknowledging the existence of internal preferred states (so called ‘reference variables’) arranged in a control hierarchy, together with the power of negative feedback to achieve and maintain those preferred states in the face of unpredictable disturbances. Crucially, PCT focuses attention on the control of perceptual *inputs* rather than behavioural *outputs* [68].

Of course, negative feedback control is endemic in the field of robotics (particularly at the lowest levels of motor control), not least because many robotocists are trained control engineers. However, mainstream control engineering tends to concentrate on the regulation of required behaviour (low-level outputs) in a single control loop, and the emphasis is on modelling the dynamics of the system under control in order to calibrate the parameters of the control process. Also, hierarchical control structures are much less common. PCT thus offers a broader perspective that is often overlooked, as well as introducing the proposal that living organisms control *perceptions* rather than behaviours.

2.1 Classic Automatic Control

As a reminder, PCT was developed in the context of existing knowledge of the classic theory of automatic control - the study of systems that are capable of self-regulation, usually through the mechanism of negative feedback [69]. A classic single-input single-output closed loop negative feedback control system from automatic control theory is illustrated in Fig. 6. The plant (or controlled system) g_2 is the process controlled by the feedback control system, and the behaviour of the plant may be subject to arbitrary disturbances d . The feedforward (control) elements g_1 generate control signals u that are applied to the plant which produces the controlled output c . The reference input r specifies the desired output of the plant. The feedback elements h map the controlled output c to the feedback signal b , and this is compared to the external reference input r and the resulting error signal e generates a control action via g_1 . With appropriate functions for g_1 , g_2 and h , the system should compensate for any disturbances and stabilise with $b \rightarrow r$ *without having to measure d*. The process is termed *negative feedback* because the comparator involves subtraction.

It can be shown that the input-output relations of the classic control system shown in Fig. 6 are given (in the frequency domain) by the equation

$$C = \left(\frac{G_2 G_1}{1 + H G_2 G_1} \right) R, \quad (1)$$

where $\frac{G_2 G_1}{1 + H G_2 G_1}$ is the closed-loop transfer function and $H G_2 G_1$ is the loop gain. If $G_2 G_1 \gg 1$ and $H \approx 1$, then $C \approx R$.

Clearly the tracking behaviour of a classic negative-feedback control system depends on the feedforward (control) elements g_1 . If the control element is slow to respond to a disturbance, then stabilisation may take too long; this is referred

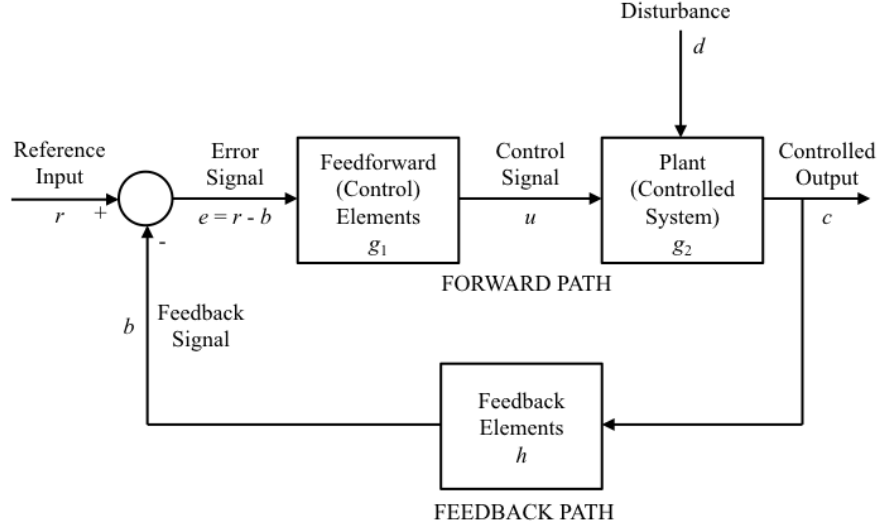


Fig. 6. Illustration of a classic negative feedback control system from automatic control theory (derived from Fig. 2-6 on page 16 of [69]).

to as an overdamped system. On the other hand, if the loop gain is too high, then the system may overshoot and even oscillate; this is referred to as an underdamped system. In practice, g_1 is often implemented using a Proportional, Integral, Derivative (PID) controller which takes the following *idealised* form:

$$u_{PID} = K_P e + K_I \int e(t) dt + K_D \frac{de}{dt}, \quad (2)$$

where the three constants - K_P , K_I and K_D - are used to optimise the stability of the control process (referred to as a critically damped system). K_P is known as the ‘loop gain’ and it determines the speed of response of the system to error. If K_P is too high, then the system may respond too quickly and overshoot (or even become unstable); if it is too low, then the system may respond too late to counteract any disturbance. The effect of K_D is to slow the rate of change of the controller output and thus minimise overshoot arising from a high K_P . K_I is intended to reduce any residual steady-state error. There are various methods for estimating these constants and, in practice, K_I and K_D can often be set to zero.

2.2 Hierarchical Perceptual Control

Contemporary automatic control systems tend to be based on a single control loop, but the dynamic behaviour of the system to be controlled is often very complex. PCT, on the other hand, decomposes a system into a hierarchical/layered structure with a multiplicity of control loops at each level; this is referred to as *Hierarchical Perceptual Control Theory* (HPCT) [65]. Inspired by classic automatic control theory, the basic perceptual control unit (feedback control loop) in PCT is illustrated in Fig. 7.

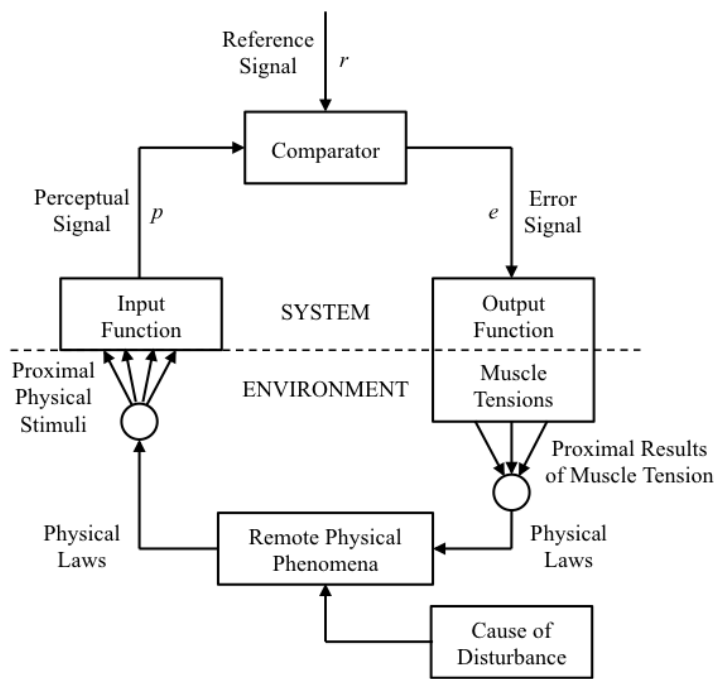


Fig. 7. Illustration of a basic perceptual control unit (derived from Fig. 5.2 on page 61 of [65]).

Note the similarity between Fig. 7 and Fig. 6. The main difference is that the PCT structure illustrated in Fig. 7 is specifically aimed at modelling the behaviour of a *living* system, and this means that the reference variable is intrinsic (rather than extrinsic). In all other respects the two control structures are effectively identical³.

³ Note that PCT practitioners refer to the feedback signal (p in Fig. 7) as the *controlled variable* thereby emphasising that it is the perceptual signal that is controlled in a negative feedback control system, not the plant.

As mentioned, HPCT derives its power from the hierarchical structure of multiple control loops in which the output from one loop provides the reference signal for another loop. Powers [65] proposed specific orders of control (starting at the lowest first-order level) as follows: *intensity* \Leftrightarrow *sensation* \Leftrightarrow *configuration* \Leftrightarrow *transitions* \Leftrightarrow *sequence* \Leftrightarrow *relationships* \Leftrightarrow *program* \Leftrightarrow *principles* \Leftrightarrow *system concepts*.

Powers also generalised the basic perceptual control unit to include the storage and retrieval of information in memory. In particular, he proposed that all reference signals constitute recordings of past perceptual signals that are retrieved by address signals from the higher level - see Fig. 8. Implementation of this feature required the addition of two switches - a memory switch and a perceptual switch. Powers noted that each of the four possible combinations of switch settings lead to interesting outcomes. If both switches are vertical ($a+d$), then the loop is in “*conventional control mode*”. If the perceptual switch is vertical and the memory switch is non-vertical ($a+c$), then information can be acquired without action taking place; Powers called this “*passive observation mode*”. If the memory switch is vertical and the perceptual switch is non-vertical ($b+d$), then control takes place with no perceptual awareness at higher levels; Powers referred to this as “*automatic mode*”. Finally, if both switches are non-vertical ($b+c$), then past perceptions (retrieved from memory) are redirected back up the hierarchy; Powers designated this as “*imagination mode*”, and identified it as a potential mechanism for visualisation and planning.

On the surface, the HPCT architecture looks remarkably similar to subsumption (as described in section 1.2). However, a key distinction is that HPCT is capable of optimising many variables simultaneously within a negative feedback control framework, whereas subsumption is essentially an event-driven *stimulus* \rightarrow *response* architecture [70].

Finally, Powers identified three types of learning within HPCT: the storage of information in memory, problem-solving and *reorganisation*. The latter is particularly interesting since it involves altering the structure and properties of the control systems themselves. Powers viewed reorganisation as a kind of meta control system that had the objective of reducing the intrinsic error of the overall system to zero.

2.3 The Way Forward?

Perceptual control theory (and HPCT) appears to offer a number of interesting features for modelling the behaviour of living systems [71], and should thus be a serious contender for implementing control in artificial systems such as robots (as envisaged by Powers [72,73,74,75]). A few attempts have been made in this area by PCT practitioners [76,77,78]. However, since a large proportion of robotocists are control engineers who take negative feedback systems as a given, they have not seen the particular benefits of using (H)PCT.

On the other hand, (H)PCT has started to have some influence in the field of spoken language processing [79]. It was established many years ago that human speech is optimised *by speakers for listeners* [80], and this has been theorised

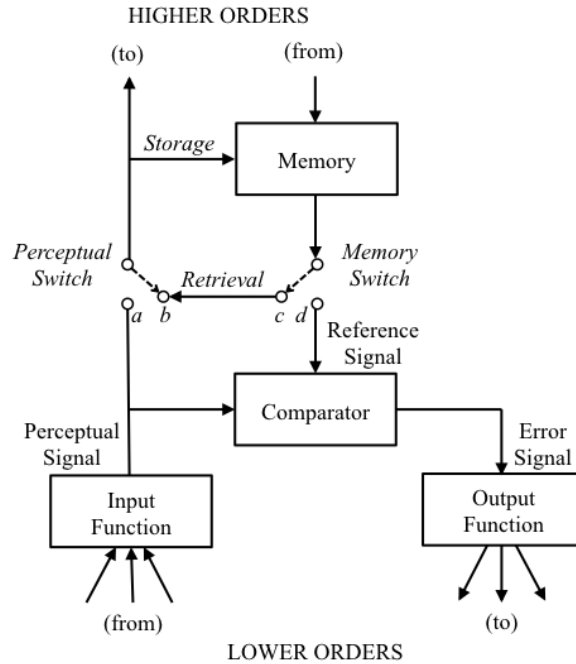


Fig. 8. Illustration of an HPCT unit showing the perceptual and memory switches which can be set for control mode ($a+d$), passive observation mode ($a+c$), automatic mode ($b+d$) or imagination mode ($b+c$) (derived from Fig. 15.2 on page 221 of [65]).

as resulting from the operation of negative feedback processes [81]. Hence, there is an obvious link with PCT. However, mainstream speech technology systems have tended to ignore speaker-listener dependencies and such behaviour is simply modelled as stochastic variation [82].

In an attempt to capture such dependencies, and inspired by classic automatic control, (H)PCT and recent discoveries in cognitive neuroscience (such as mirror neurons - discussed in Section 1.3), the PREdictive SENSory Control and Emulation (PRESENCE) architecture [83] provides some novel technical solutions in this area. For example, a PRESENCE-based text-to-speech synthesiser - *C2H* - has been developed that is capable of adjusting its pronunciation while it is speaking as a function of its perceived communicative success, the latter being judged using a perceptual feedback path in which the ‘controlled variable’ was an estimate of the intelligibility [84,85].

PRESENCE not only provides a novel architecture for generating spoken language, it also shows how the principles of feedback control can be applied to the recognition and interpretation of spoken language [86]. Indeed, PRESENCE emphasises that in order to manage its own behaviour (its perceptions), a social agent also needs to be able to interpret the world and the behaviour of other in-

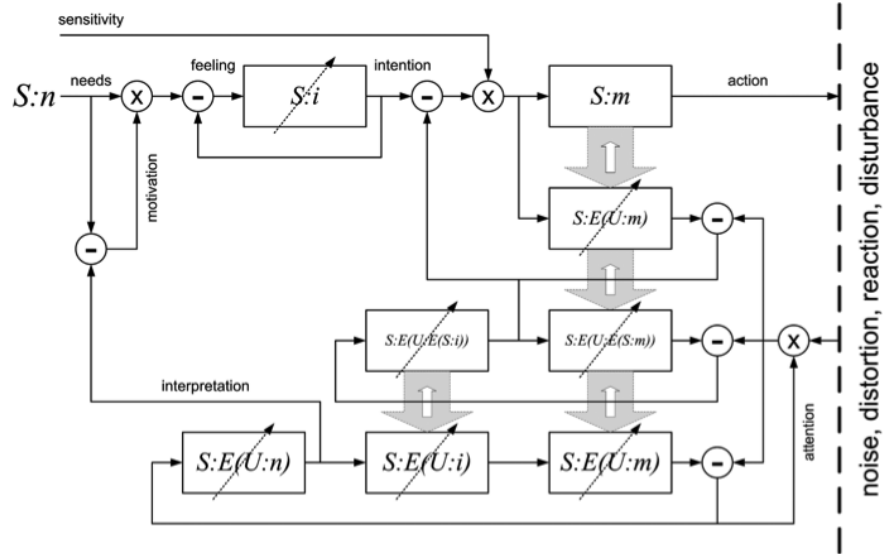


Fig. 9. Illustration of the PCT-inspired PREdictive SENsorimotor Control and Emulation (PRESENCE) architecture for spoken language processing [79,83]). The diagram depicts a high-level view of an integrated speaker-listener; the first level is the primary route for motor behaviour (speaking), the second layer is an emulation of possible motor actions, the third layer is an emulation of the interlocutor’s emulation of the speaker-listener and the bottom layer is the interpretation of the interlocutor (listening). All levels involve control feedback processes, and information is shared between one layer and another. The architecture is intended to describe a system that can communicate successfully in the face of disturbances such as noise, distortion and feedback from the interlocutor.

telligent systems within it. Hence, much greater emphasis is placed on modelling perception than is the norm in mainstream PCT. PRESENCE thus provides an interesting insight into a potentially more general-purpose PCT-inspired architecture that is not specific to spoken language processing and which could be applied in engineered solutions. However, the founding principles of such an architecture have not been fully elucidated, so what follows is a step in that direction.

3 Towards ‘Intelligent’ Communicative Systems

The insightful catchphrase for (H)PCT is “*Behaviour is the control of perception*” - but what is perception, and how are controlled variables estimated? In

B:CP [65], Powers himself does not say a great deal about perceptual mechanisms; he acknowledges the potentially arbitrary relationship that exists between internal perceptions and external reality, but he seems to assume some form of straightforward (neural) transformation/mapping from first-order intensities to *invariant* higher-order perceptions that is carried out by an ‘input function’ (albeit mediated by memory - see Section 2.2). Of course such an assumption is common in the pattern recognition/machine learning literature. However, in practical fields such as visual or auditory scene analysis (and even in the latest theories of brain function [87]), the notion of ‘perception as transformation’ has been complemented by the emergence of more powerful (and more successful) *generative model-based* solutions [88,89,90]. Hence, there is a potentially valuable opportunity to map these findings back into (H)PCT.

Not only is this an interesting route to take, but it also turns out to be a crucial step towards a deeper understanding of how one intelligent system (such as an autonomous social robot) might interact/communicate with another (such as a human being) [91]. What follows, therefore, is a fundamental analysis - starting from first principles - that intersects with (H)PCT and then extends it in important and interesting ways.

3.1 Actions and Consequences

Consider a world that obeys the ordinary Laws of Physics. The world W has a set of possible states S , and its state $s[t]$ at time t is some function of its state $s[t-1]$ at time $t-1$. The world can thus be viewed as a form of dynamical system (probably non-linear, almost certainly stochastic) that evolves from state to state as time progresses. These state transitions can be expressed as a *transform* ...

$$f_W : s[t-1] \rightarrow s[t], \quad (3)$$

where f_W is some function that transforms the state of the world at time $t-1$ to the state of the world at time t .

This means that the evolution of events in the world constitutes a continuous cycle of cause-and-effect. Events follows a time course in which it can be said that *actions* (i.e. the sequence of events in the past) lead to *consequences* (i.e. the sequence of events in the future) which constitute further *actions*, leading to further *consequences*, and so on ...

$$\text{Consequences} = f_W(\text{Actions}). \quad (4)$$

This straightforward scenario (as described by equations 3 and 4) can be expressed diagrammatically as shown in Fig. 10.

Of course, while the transform f_W might be relatively simple (since it is based on the Laws of Physics), the state-space S of possible *actions* and *consequences* could be immense depending on the complexity of the world W . This means that it is impossible to model everything that happens in the world. However, in practice, some parts of the world might have very little influence on other parts. So it is possible to consider a subset of the world w that has a minimal dependency on the rest of the world \bar{w} .




Fig. 10. Illustration of the continuous cycle of cause-and-effect in a world that obeys the ordinary Laws of Physics.

3.2 An Agent Manipulating the World

Now consider the presence of an intentional agent a (natural or artificial) that seeks to effect a change in the world⁴. In this case the agent's intentions are converted into actions which are, in turn, transformed into consequences ...

$$\text{Consequences} = f_w(g_a(\text{Intentions})), \quad (5)$$

where g is some function that transforms the agent's intentions into actions (a process known in robotics as 'action selection') - see Fig. 11.




Fig. 11. Illustration of an intentional agent manipulating the world in an open-loop S-R configuration.

The situation illustrated in Fig. 11 corresponds to an open-loop S-R configuration, and this means that the accuracy with which an agent can achieve its intended consequences is critically dependent on having precise information about both f and g . Mathematically, the best method for achieving the required consequences is for the agent to employ an *inverse transform* in which g is replaced by f^{-1} - see Fig. 12⁵.

⁴ The reason *why* the agent wishes to change the state of the world is addressed later (in Section 3.7).

⁵ Note that, in the situation where f represents the plant operating a robot's actuators, the estimation of f^{-1} is commonly referred to as 'inverse kinematics'.




Fig. 12. Illustration of an intentional agent manipulating the world in an open-loop S-R configuration using an inverse model.

It is possible to discuss at length how information about the transforms g , f or f^{-1} could be acquired. For example, model parameters could be computed using complex mathematical tools for system estimation or by employing machine learning techniques on extensive quantities of training data using a process known as ‘expectation maximisation’ (EM) [92]. Whichever approach is taken, the final outcome would inevitably be sensitive to any inaccuracies in calibrating the relevant model parameters as well as being unable to tolerate unforeseen noise and/or disturbances present in the agent or in the world.

Control theory, and thereby PCT, of course provides an alternative *closed-loop* solution that is not dependent on knowing f (or f^{-1}). An agent simply needs to be able to judge whether its objectives are being met - that is, whether the consequences of its actions *match* its intentions. An agent thus needs to be able to choose actions that minimise the difference between its intentions and the perceived consequences of its actions - see Fig. 13. In PCT terminology, the agent’s intentions correspond to the reference signal and the perceived consequences correspond to the controlled variable; actions are selected in order to minimise the difference between the two.




Fig. 13. Illustration of an intentional agent manipulating the world in a *closed-loop* PCT-style negative-feedback configuration. In this configuration the transform g corresponds to the ‘output function’ (in PCT) or ‘controller’ (in classic control theory), and the transform h corresponds to the ‘input function’ (in PCT) or ‘feedback element’ (in classic control theory).

If the perceptual signal, and hence the error signal, is in the same parameter space as the control signal (see Fig. 6), then it is possible - in principle - to jump to the optimum solution in one step. However, in practice it takes time to correct an error (since physical actions cannot take place instantaneously). So the process typically *iterates* towards a solution (by information flowing around the loop). If, on the other hand, the perceptual signal and the control signal are not in the same parameter space, then there is no direct mapping between error and control action (unless such a mapping has been learnt in advance). So, again, an iterative procedure is required⁶.

This means that, although closed-loop control does not require information about f (or f^{-1}), it does need to know about g - the relationship between the error gradient in perceptual space and the appropriate control action. Even a simple thermostat needs to be configured correctly such that a decrease in the temperature of a room leads to an increase in the level of heating (and *vice versa*); in this case, the uni-dimensional perceptual error gradient determines the polarity of the action of the controller. In general, perceptions and actions may lie in (different) high-dimensional spaces⁷ and, either the relationship between possible behaviours and the perceptual error gradient is (somehow) known in advance, or it has to be discovered (learnt) by active exploration; for example, using ‘reinforcement learning’ [93] or the process referred to in PCT as ‘reorganisation’ - see Section 2.2.

A further complication is that the error space may not be continuous. In this case there is no gradient and hence no indication of which action to select in order to arrive at the correct solution. For example, imagine a person who wishes to illuminate one section of a large room being faced with a set of unmarked light switches laid out in no obvious arrangement. The only available strategy is to try each switch in a random fashion until the appropriate light illuminates⁸. Of course, once a mapping - g - has been discovered, it may be stored in memory and used to guide future behaviour.

In many situations, negative feedback control is able to employ an optimisation technique known as *gradient descent* in which the difference between the intentions (the reference signal) and the perceived consequences (the feedback signal) is a continuous variable that can be reduced monotonically to zero. In this case, the dynamics of the optimisation is dependent on the appropriate choice of parameters in the controller (for example, K_P , K_I and K_D in a classic PID controller - as described in Section 2.1); an inappropriate choice of controller

⁶ For example, a thermostat measures temperature, whereas the control signal for a heating/cooling system would be expressed in terms of power output. This means that *a-priori* there is no information on how to convert a difference between desired and actual temperature (the error signal) to an optimum power setting (the control signal). Hence the power output has to be adjusted *iteratively* until the desired temperature is reached.

⁷ In robotics, the space of possible actions is characterised by the Degrees-of-Freedom (DoF) in the system, and this usually corresponds to the number of joints or actuators/motors.

⁸ This is a common experience for an academic teaching in an unfamiliar lecture room.

parameters (such as the wrong sign for K_P) can even send the system in the opposite direction (that is, *positive feedback*) with a consequent failure to converge. In classic control systems, the process of optimising control parameters is often performed manually. In PCT, it is performed by reorganisation (see Section 2.2) which is, itself, a negative feedback control process (and which is interestingly similar to the ‘actor-critic’ approach to reinforcement learning [94]).

This discussion makes it clear that, in general, **negative feedback control may be viewed as an iterative search⁹ process** and, in the example of the discrete light switches with no prior information, only a *random* (as opposed to *directed*) search is feasible. The realisation that control can be regarded as a type of search is a key result, and it emphasises that the topology of the search space and the controller’s information about that topology is key to the effectiveness and outcome of the optimisation.

Returning to the agent a attempting to manipulate the world w , negative feedback control can thus be viewed as a search over possible actions to find those which give rise to the best match between intentions and perceived consequences. This can be expressed as

$$\widehat{Actions} = \arg \min_{Actions} (Intentions - Perceived Consequences), \quad (6)$$

where $\widehat{Actions}$ represents an estimate of the actions required to minimise the difference between intentions and perceived consequences - see Fig. 14.




Fig. 14. Illustration of an intentional agent manipulating the world in a *closed-loop* configuration using a search process to select actions that minimise the difference between its intentions and the perceived consequences.

The structure illustrated in Fig. 14 is similar to the classic PCT arrangement, but enhanced by the identification of iterative *search* as the basis of a negative-feedback controlled optimisation process. However, this configuration will only function correctly if (i) the agent can observe the consequences of its actions

⁹ Note that ‘search’ is a term used by computer scientists to describe a process that engineers would refer to as ‘optimisation’. In both cases, a system iterates/converges towards an optimum solution by minimising/maximising some ‘objective function’ - in this case, minimising an error signal. The value of using the term ‘search’ is that it covers both continuous and discrete forms of optimisation.

and (ii) the search space is concave¹⁰. If the consequences of an agent's actions are hidden from observation in *space* (because objects are obscured¹¹) or in *time* (because feedback delays are too high or the intended consequences are in the future), the loop can still function, but only if the agent is able to *estimate* the consequences of possible actions. Likewise, if the search space is not concave but has many local minima, then an iterative search can avoid getting stuck in the nearest one by exploring the whole space *in advance*. In other words, in both of these cases, an agent could gain benefit by being able to *predict* the consequences of possible actions.

In other words, if an agent cannot directly observe the consequences of its actions, or the search space has many local optima, then it needs to (i) estimate the relationship between possible actions and possible consequences (f_w), (ii) perform a search over hypothetical actions and then (iii) execute those actions that are found to minimise the estimated error. In this case ...

$$\widehat{Actions} = \underset{\widehat{Actions}}{\operatorname{arg\,min}} \left(Intentions - \widehat{f}_w(\widehat{Actions}) \right), \quad (7)$$

where \widehat{f}_w is the estimate of f_w and $\widehat{Actions}$ is the set of possible actions - see Fig. 15.




Fig. 15. Illustration of an intentional agent manipulating the world in a *closed-loop* PCT-style configuration in the situation where it is unable to directly observe the consequences of its actions (in space or time).

What is interesting in this arrangement is that the estimated transform \widehat{f}_w can be interpreted as a form of mental simulation (or predictor) which emulates the consequences of possible actions prior to action selection [95]. In other words, searching over $\widehat{f}_w(\widehat{Actions})$ is equivalent to *planning* (and corresponds to Powers' *imagination mode* in HPCT, i.e. the *b+c* setting in Fig. 8). It also shows that the ability to perform “what if” simulations is not a property of a particular

¹⁰ That is, the search space has only one *global* optimum.

¹¹ For example, attempting to achieve some objective in the dark, reaching behind another object or, of particular significance here, manipulating the internal states of another agent (as will be discussed in Section 3.4).

level of control, but is a generic property that can be instantiated in any control loop at any level (as proposed by Powers).

The ability to emulate the consequences of possible actions is important because, not only is it a link between PCT and AI planning, but it also offers the possibility of finding global (rather than local) solutions¹². From an ecological perspective, this could be critical in terms of saving living systems vital time and/or energy in real world situations.

3.3 An Agent Interpreting the World

Having established an extended PCT-style framework for an agent attempting to manipulate the physical world (with search as the underlying mechanism supporting negative feedback control), it is now possible to turn to the complementary situation in which an agent a is attempting to *interpret* the world w . In this case, interpretation is defined as an agent deriving potentially hidden actions/causes of events by observing their visible effects/consequences¹³.

$$\widehat{Actions} = h_a(f_w(Actions)), \quad (8)$$

where h is some perceptual function that transforms observed effects into estimated causes - see Fig. 16.




Fig. 16. Illustration of an agent a attempting to infer the hidden causes/actions from their observable effects/consequences.

Given that consequences are caused by actions via the transform f_w , it is possible (in principle) to compute the actions directly from the observed consequences using the inverse transform f_w^{-1} - see Fig. 17.

Of course, the accuracy of this process depends on the fidelity of the inverse transform. In practice, f_w^{-1} is not known and very hard to estimate. A more

¹² As an example, the fastest route between two points on a map may not be the shortest, nor does each move forward necessarily take you nearer to the final destination. This is the difference between local and global optimisation.

¹³ For example, given an observation of an object suddenly accelerating along a flat surface, it is possible to *infer* (using the laws of physics) that a hidden force must have acted upon that object.




Fig. 17. Illustration of an agent a attempting to infer the hidden causes/actions from their observable effects/consequences by means of the inverse transform f_w^{-1} .

tractable solution is thus not to use an inverse model, but to construct a so-called ‘forward model’¹⁴ (that is, an estimate of f_w) and to compare its output with the observed signals. Mathematically, this is the optimum way to estimate hidden variables¹⁵ given uncertainty in both the observations and the underlying process. Such a configuration is known as a ‘maximum likelihood (or *Bayesian*) classifier’. Also, it is a standard result in the field of statistical estimation that the parameters of forward/generative models are typically much easier to derive using reliable techniques for maximum likelihood (ML) or maximum *a-posteriori* (MAP) estimation, and performance degrades gracefully when there is missing data.

The process of interpretation using such an arrangement thus proceeds by searching over possible actions/causes to find the best match between the predicted and the observed consequences ...

$$\widehat{Actions} = \arg \min_{Actions} (Consequences - \widehat{f}_w(Actions)). \quad (9)$$

This process is illustrated in Fig. 18, and what is immediately apparent is that, just as the manipulation case, the process of **interpretation may also be construed as a negative-feedback control loop** (in this case, a ‘search’ over possible explanations). This is a significant outcome that is not explicit in (H)PCT; hence it provides a potentially valuable link between PCT and mainstream approaches to machine perception.

In fact, the architecture illustrated in Fig. 18 is a familiar¹⁶ model-based *recognition* framework in which the recognition/interpretation/inference of the (hidden) cause of observed behaviour is viewed as a search over possible outputs from a forward model that is capable of generating that behaviour [96,97]. This is an established and powerful approach in machine perception and scene analysis

¹⁴ Also known as a ‘generative model’.

¹⁵ A ‘hidden variable’ is a measurement which cannot be made directly (e.g. by observation).

¹⁶ ... in the field of machine learning.




Fig. 18. Illustration of an agent a attempting to infer the hidden causes/actions from their observable effects/consequences by using a negative-feedback control loop to search the outputs from $\widehat{f_w}$ (a forward estimate of f_w).

(known as *analysis-by-synthesis*¹⁷), and it has been applied very successfully in a range of practical scenarios [98,88,99].

3.4 An Agent Communicating its Intentions to Another Agent

The foregoing establishes a remarkably symmetric framework for manipulating and interpreting the world in the presence of uncertainty and unknown disturbances; both employ PCT-style negative-feedback control loops that perform a search over the potential outputs of forward models. This section (and the next) extends the arguments to the case where the world contains more than one agent: a world in which a sending agent s is attempting to change the mental state of a receiving agent r (that is, *communicating* its intentions without being able to directly observe whether those intentions have been perceived¹⁸) and the receiving agent is attempting to interpret the sending agent (that is, estimating the sending agent's intentions based only on observations of its actions).

So, again starting from first principles - for the sending agent s

$$Actions_s = g_s(Intentions_s), \quad (10)$$

where g_s is the transform from intentions to behaviour, and for the receiving agent r

$$Interpretations_r = h_r(Actions_s), \quad (11)$$

where h_r is the transform from observed behaviour to interpretations - see Fig. 19.

Hence, for agent s attempting to communicate its intentions to agent r , the arguments put forward in Section 3.2 suggest that, if there is no direct feedback

¹⁷ The term 'analysis-by-synthesis' refers to an established process in the field of signal processing in which the value of an unknown variable is estimated from observation data ('analysis') using a model of how the observations might have been generated ('synthesis').

¹⁸ It might seem perverse to rule out feedback. However, the aim here is to consider the general case, where feedback may be present or absent.




Fig. 19. Illustration of a world containing two communicating agents - a sender s and a receiver r .

from agent r , then agent s needs to compute appropriate behaviour (actions) based on

$$\widehat{Actions}_s = \arg \min_{\widehat{Actions}_s} \left(Intentions_s - \widehat{h}_r(\widehat{Actions}_s) \right), \quad (12)$$

which is a negative-feedback control loop performing a search over possible behaviours by agent s and their interpretations by agent r as estimated by agent s - see Fig. 20. This process can be viewed as *synthesis-by-analysis*.




Fig. 20. Illustration of a sending agent optimising its communicative behaviour by inferring the interpretations made by the receiving agent using a process of *synthesis-by-analysis*.

3.5 An Agent Interpreting the Behaviour of Another Agent

For agent r attempting to interpret the intentions of agent s , the arguments put forward in Section 3.3 suggest that agent r needs to compare the observed actions of agent s with the output of a forward model for agent s

$$\widehat{Intentions}_s = \arg \min_{Intentions_s} \left(Actions_s - \widehat{g}_s(Intentions_s) \right). \quad (13)$$

which is a negative-feedback control loop performing a search over the possible intentions of agent s and their realisations by agent s as estimated by agent r - see Fig. 21. As in Fig. 18, this process can be viewed as *analysis-by-synthesis*.




Fig. 21. Illustration of a receiving agent inferring the communicative intentions of a sending agent using a process of *analysis-by-synthesis*.

Interestingly, this configuration is exactly how contemporary approaches to automatic speech recognition are formulated (using a probabilistic forward generative model known as a ‘Hidden Markov Model’ - HMM) [90]). In fact, the analysis-by-synthesis approach to speech recognition is not only reminiscent of the *Motor Theory* of speech perception [100], but is also supported by recent neuroimaging data [101,102]. Hence, there is ample evidence that such a configuration is appropriate for modelling perceptual inference.

3.6 Using *Self* to Model *Other*

Looking at the arrangements outlined in Sections 3.4 and 3.5, we arrive at an important result; both require one agent to have a model of (some aspect of) the other. The sending agent s selects its actions by searching over possible interpretations by the receiving agent r using an estimate of the receiving agent’s transform from observations to interpretation (\widehat{h}_r). The receiving agent r infers the intentions of the sending agent s by searching over possible interpretations using as estimate of the sending agent’s transform from intentions to actions (\widehat{g}_s) - see Fig. 22.




Fig. 22. Illustration of a world containing one agent (a sender s) communicating with another (a receiver r) where each makes use of a model of the other. The sender infers the interpretation of its intentions by the receiver using *synthesis-by-analysis*, and the receiver infers the intentions of the sender using *analysis-by-synthesis*.

The configuration shown in Fig. 22 leads to an interesting question: where do the transforms \widehat{h}_r and \widehat{g}_s come from? More precisely, how might their parameters be estimated? Obviously they could be derived using a variety of different learning procedures (including PCT-style reorganisation). However, one intriguing possibility is that, because of the similarity between agents¹⁹, **each agent could approximate these functions using information recruited from their own structures**²⁰. In other words, $\widehat{h}_r \leftarrow h_s$ (which can be searched using g_s rather than \widehat{g}_r) and $\widehat{g}_s \leftarrow g_r$ (which can be searched using h_r rather than \widehat{h}_s) - see Fig. 23.

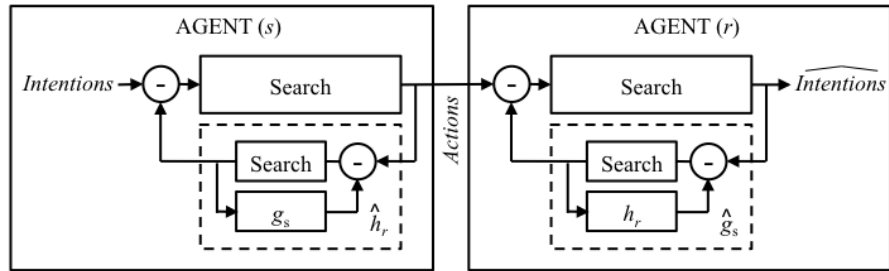


Fig. 23. Illustration of sending and receiving agents exploiting knowledge of themselves in order to communicate.

This arrangement, in which both agents exploit sensorimotor knowledge of themselves to model each other, can be thought of as *synthesis-by-analysis-by-synthesis* for the sending agent and *analysis-by-synthesis-by-analysis* for the receiving agent. Combining both into a single communicative agent gives rise to a structure where perception and production are construed as parallel recursive control feedback processes (both of which employ search as the underlying mechanism for optimisation), and in which the intentions of self and the intentions of other are linked to the behaviour of self and the observations of other, respectively - see Fig. 24.

3.7 A Needs-Driven Communicative Agent

The preceding arguments have provided interesting answers to two key questions: (i) how can an agent optimise its behaviour in order to communicate

¹⁹ This situation is particularly relevant to the high degree of similarity that is found in living systems between conspecifics (members of the same species).

²⁰ The idea of recruiting information about *self* in order to model *other* is a core component of the PRESENCE architecture (Section 2.3) which, in turn, is founded on the principle of mirror structures in the brain (Section 1.3).

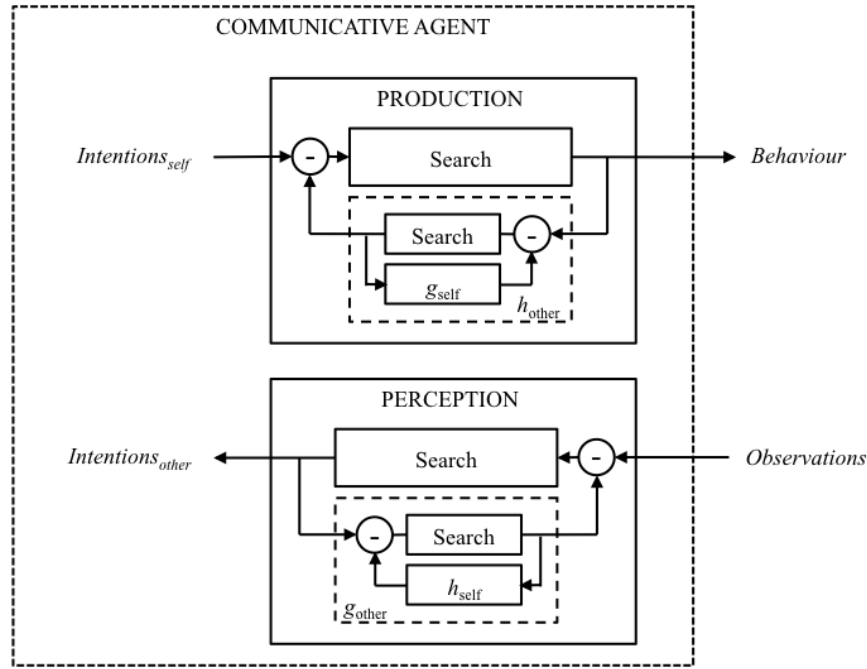


Fig. 24. Illustration of a communicative agent that is capable of optimising the signalling of its own intentions (PRODUCTION) and inferring the intentions of others (PERCEPTION).

its intentions and (ii) how can an agent infer the intentions of another agent by observing their behaviour? However, thus far it has been assumed that intentionality is a key driver of communicative interaction. Hence, an obvious question is - where do the intentions come from? Of course, HPCT provides a clear answer; the reference signals at one level in a hierarchy are set by the level above and mediated by memory (as illustrated in Fig. 8). So this would suggest that intentions should be the output from some higher level.

In B:CP [65], Powers doesn't use the term 'intention' as such, rather he talks about purposeful behaviour (at *all* levels of the HPCT hierarchy). However, by invoking intentionality as a manifestation of purposeful goal-driven behaviour, it is possible to make a direct link between HPCT and contemporary architectures such as DAC (Section 1.3) and BDI (Section 1.4)²¹. In particular, the BDI approach to agent-based modelling makes it clear that intentionality is not only derived from an agent's longer-term goals and desires, but that it is also

²¹ The main difference between HPCT and these architectures (including the one developed here) is that HPCT provides a rich decomposition of purposeful behaviour into the necessary levels of detailed control, whereas the other schemes tend to collapse such levels into a single abstraction.

conditioned by the beliefs that an agent holds. More directly, the DAC architecture emphasises that behaviours are ultimately driven by a motivational system based on an agent's fundamental *needs* (such as energy for survival). This appears to be similar to the important role that intrinsic control systems play in the development of the perceptual hierarchy in HPCT [65].

Putting all this together, it is possible to formulate a generic (and remarkably symmetric) architecture for a *needs-driven* communicative agent that is both a sender and a receiver²² - see Fig. 25. In this framework it is proposed that a communicative agent's behaviour is conditioned on appropriate motivational and deliberative states (reference signals): *Needs* \rightarrow *Desires* \rightarrow *Intentions*. Likewise, the intentions, desires and needs of another agent are inferred via a parallel interpretive structure: *Perception* \rightarrow *Interpretation* \rightarrow *Comprehension*²³.

4 Discussion

The foregoing section has established a putative framework for a needs-driven communicative agent in which the data structures for needs, desires and intentions represent the informational *belief* states (controlled variables) of the agent. As such, these representational structures are analogous to the memory layers in HPCT (as shown in Fig. 8). However, a key feature of the configuration proposed here (for a communicative agent) is that these structures explicitly contain information relating to both self and other, and this allows the behaviour of self to be conditioned on the inferred informational state of other. In other words, it provides a basis for modelling interpersonal stances such as *empathy* [104] and *Theory of Mind* (ToM) [105]²⁴.

Another interesting aspect of the architecture illustrated in Fig. 25 is that, like DAC and PRESENCE, it is founded on a motivational system based on *needs* (for example, as described by Maslow in his famous 'Hierarchy of Needs' [107]). This might seem to be an unnecessary embellishment - a feature that is more relevant to modelling a living system than to designing an artificial agent. However, invoking such a needs-based framework answers a fundamental question - why would an agent (natural or artificial) do anything [108]? Also, it is envisaged that it would be useful to partition needs into those that pertain to physical health (such as survival and safety needs) and those that pertain to psychological health (such as personal and social needs), with the relevant feedback signals being related to physiological and mental well-being. This not

²² A forerunner of this framework has been referred to as MBDIAC (Mutual Beliefs Desires Intentions Actions & Consequences) [103].

²³ From a PCT perspective, these are all perceptions.

²⁴ Interestingly, Marken - a long-standing PCT practitioner - has recently posited the value of PCT in testing the validity of the inferences that observers make about the intentional states of others [106]. This might be usefully carried over into the proposed framework as a strategy an agent might employ in order to establish a desired cooperative/competitive relationship with another agent.

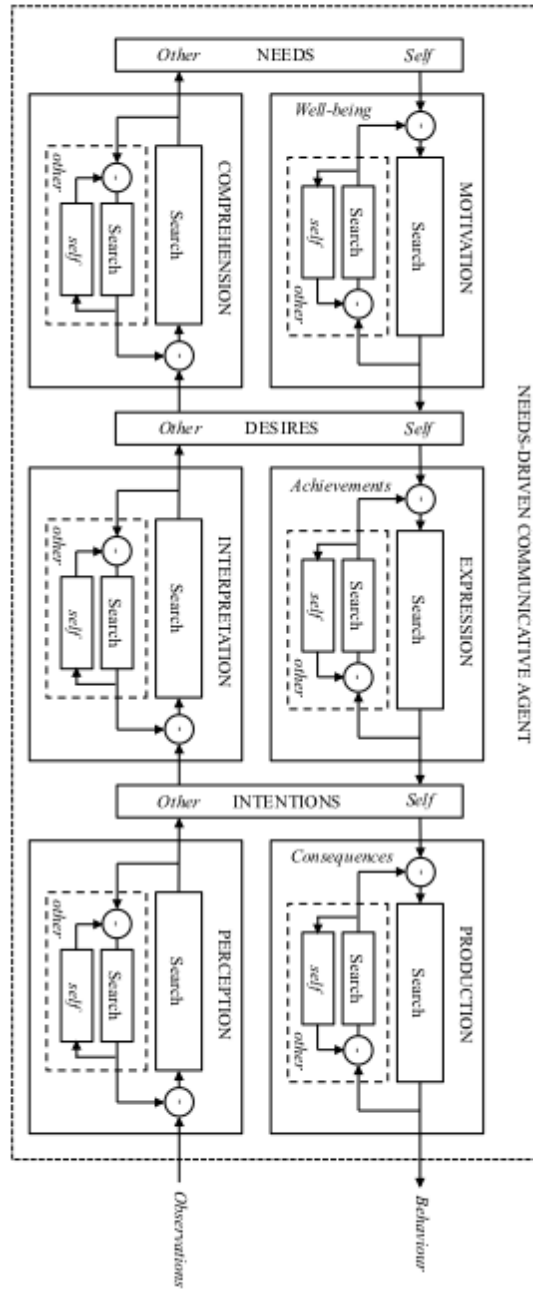


Fig. 25. Illustration of the proposed architecture for a needs-driven communicative agent.

only ties in closely with Powers' observations concerning the importance of optimising *intrinsic* variables that are related to the health of the organism, but it also links with his view that the control of such variables provides the basis for reorganisation (i.e. structural learning) [65].

In practice, the motivational framework can be decomposed into two elements: (i) a needs structure that provides the incentives for behaviour and (ii) the amount of effort that will be devoted to meeting those needs. The previous paragraph (and the architecture shown in Fig. 25) addresses the first of these. The second relates to the *enthusiasm* of an agent, i.e. how much it cares about meeting its needs. In a conventional PCT-style negative-feedback control loop, such behaviour would correspond to the loop gain; a high loop gain implying a high degree of physical effort/enthusiasm and *vice versa*. In the framework outlined here, such behaviour would relate to the depth and quality of the search(es) that pervade the overall structure - deep search(es) implying a high degree of mental effort/enthusiasm and *vice versa*.

Interestingly, the proposed framework also provides a practical architecture that supports appraisal theories of emotion [109,110] based on the classic dimensions of pleasure and arousal [111]. For example, the comparators that pervade the structure each constitute mini-appraisal units that generate measures of positive/negative affect - the error signals [112]. Likewise, the effort that an agent devotes to achieving its goals and intentions (as discussed in the previous paragraph) can be interpreted as a measure of arousal. Tying this in with Powers' notion of intrinsic variables driving reorganisation, these emergent properties of the proposed framework can be interpreted as (i) emotion driving behaviour and (ii) the perception of emotion driving learning²⁵.

Finally, the core components in the proposed framework have been portrayed in Fig. 25 as a three-layered control structure. However, this is just a necessary simplification in order to establish the basic principles. Also, the emphasis has been on a communicative agent that seeks to influence the internal states of another agent (as opposed to an agent that seeks to manipulate the physical behaviour of another agent). In practice, the structures outlined here would be decomposed into a richer hierarchy of parallel control loops - much as envisaged in HPCT. The main difference would be that the perceptual apparatus in the proposed framework would be more structured than in HPCT, with explicit parameter sharing between behavioural and perceptual components (rather than the switching arrangement illustrated in Fig. 8). Also, whilst Powers' mechanism for reorganisation should be capable of optimising the structure of a decomposed architecture for any given problem, in practice no-one has yet succeeded in developing a practical solution. Since some of the solutions presented here are related to more contemporary approaches to machine learning, it is envisaged that they might offer a way forward in understanding how to configure an HPCT structure automatically.

²⁵ This overlaps nicely with Powers' description of the role of emotion [65].

5 Summary and Conclusion

This chapter has shown how a practical framework for modelling and implementing ‘intelligent’ systems can be developed from the establishment of a set of fundamental principles which both intersects with and extends Perceptual Control Theory. It has been argued that PCT has hitherto placed less emphasis on the transformation of perceptual functions through mental simulation, and has treated the mapping from sensation to perception as a relatively straightforward series of transformations (albeit mediated by memory), thereby overlooking the potential of powerful generative model-based solutions that have emerged in practical fields such as visual or auditory scene analysis. Starting from first principles, and considering how an intelligent agent might interact with the world and with other agents it might contain, it has not only been shown how these ideas might be integrated into PCT, but also how PCT might be extended towards a remarkably symmetric architecture for a needs-driven communicative agent.

The arguments presented in Section 3 first led to the realisation that the optimisation which takes place within a negative feedback control loop is equivalent to a search process, and this led to a reformulation of intentional behaviour as a search over possible actions in order to find those which minimise the difference between an agent’s intentions and the perceived consequences. These principles were then extended to cover the situation in which the consequences of actions might be hidden from direct observation (for example, when one agent is manipulating the mental state of another agent), and this led to the emergence of simulation as a key mechanism for maintaining a control loop structure by using an estimate/prediction of possible action consequences. It was then revealed how searching a simulation facilitates off-line planning which has the important benefit of being able to compute global (rather than local) solutions.

These principles were then extended to an agent attempting to interpret behaviour, and it was shown how the perceptual process itself can be configured using a familiar PCT-style negative feedback control loop (also involving simulation) which turns out to be equivalent to established techniques for model-based recognition. It was then shown how the arguments could be extended to cover agent-to-agent communication, and this led to the crucial realisation that the parameters of the models used by one agent for simulating another agent could be derived from the models used to manage their own behaviour - that is, using *self* to model *other*. This novel arrangement was then embedded within a BDI (beliefs, desires, intentions) structure to create a new framework for a needs-based communicative agent.

One of the encouraging outcomes of this analysis is that the new perspective not only relies on a central role being played by PCT-style negative-feedback control in action selection *and* interpretation, but it also links with recent discoveries in cognitive neuroscience (such as the sensorimotor overlap attributed to mirror neurons) and aligns in interesting ways with neurally-inspired architectures such as DAC (Section 1.3). Also, since it has been shown that negative feedback control (as an optimisation process) may be viewed as a type of search, this opens

up a potentially important link with a range of model-based techniques that are already established in some areas of contemporary AI and Cognitive Systems (particularly in the field of spoken language processing). This means that the development of the approach can benefit from progress being made in these other fields, as well as extending the scope and influence of PCT to provide a more parsimonious computational framework for ‘intelligent’ systems (whether natural or artificial).

Based on the framework established thus far, the critical next steps are to formulate practical mechanisms for learning local/global structures and to investigate the implications of simulations of self and other using small scale agents (such as robots) in tasks which involve real-world interaction and communication with other agents. It is too soon to extend the principles to complex scenarios (such as full blown language-based interaction). Rather, it is necessary to scale-up the complexity in a careful and controlled manner in order to maintain a firm scientific and mathematical foundation for the entire enterprise. Such work is already underway, and early results are encouraging.

Finally, the arguments presented herein lead to the following overall conclusion with respect to modelling intelligent communicative agents - if behaviour is the control of perception (the central tenet of Perceptual Control Theory), then perception (at least for communicative agents) can be said to be the *simulation* of behaviour.

6 Acknowledgements

The author would like to thank colleagues in the Sheffield Centre for Robotics (SCentRo) and the Bristol Robotics Laboratory (BRL) for discussions relating to the content of this chapter. This work was partially supported by the European Commission [grant numbers EU-FP6-507422, EU-FP6-034434, EU-FP7-231868, FP7-ICT-2013-10-611971] and the UK Engineering and Physical Sciences Research Council [grant number EP/I013512/1].

References

1. Kurzweil, R. (1990). *The Age of Intelligent Machines*. MIT Press.
2. Roy, N., Baltus, G., Fox, D., Gemperle, F., Goetz, J., Hirsch, T., Margaritis, D., Montemerlo, M., Pineau, J., Schulte, J., & Thrun, S. (2000). Towards Personal Service Robots for the Elderly. In *Workshop on Interactive Robots and Entertainment* (WIRE 2000).
3. Kanda, T., Hirano, T., Eaton, D., & Ishiguro, H. (2004). Interactive robots as social partners and peer tutors for children: A field trial. *Human-Computer Interaction*, 19(1–2), 61–84.
4. Gockley, R., Bruce, A., Forlizzi, J., Michalowski, M. P., & Mundell, A. (2005). Designing robots for long-term social interaction. *IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS 2005). Edmonton, Alberta, Canada.

5. Levy, D. (2007). *Love and Sex with Robots, The Evolution of Human-Robot Relationships*. Harper-Collins.
6. Diehl, J. J., M.Schmitt, L., Villano, M., & Crowell, C. R. (2012). The clinical use of robots for individuals with Autism Spectrum Disorders: a critical review. *Research in Autism Spectrum Disorders*, 6(1), 249-262.
7. Winfield, A. (2012). *Robotics: A Very Short Introduction*. Oxford University Press.
8. Delcomyn, F. (2007). Biologically inspired robots. In M. K. Habib (Ed.), *Bioinspiration and Robotics: Walking and Climbing Robots* (p. 544). Vienna, Austria: I-Tech.
9. Floreano, D., & Mattiussi, C. (2008). *Bio-Inspired Artificial Intelligence: Theories, Methods, and Technologies*. MIT Press.
10. Russell, S., & Norvig, P. (2003). *Artificial Intelligence: A Modern Approach*. Prentice Hall.
11. Brunswik, E. (1956). *Perception and the Representative Design of Psychological Experiments*. Berkely: University of California Press.
12. Haugeland, J. (1985). *Artificial Intelligence: The Very Idea*. MIT Press.
13. Brooks, R. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1), 14-23.
14. Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47, 139-159.
15. Braitenberg, V. (1984). *Vehicles: Experiments in Synthetic Psychology*. Cambridge, MA: MIT Press.
16. Arkin, R. C. (1998). *Behavior-Based Robotics*. Cambridge, MA: The MIT Press.
17. Vernon, D., Metta, G., & Sandini, G. (2007). A survey of artificial cognitive systems: implications for the autonomous development of mental capabilities in computational agents. *IEEE Transactions on Evolutionary Computation, Special Issue on Autonomous Mental Development*, 11(2), 151-180.
18. Vernon, D., & Furlong, D. (2007). Philosophical foundations of enactive AI. In M. Lungarella et. al. (Ed.), *50 years of AI Festschrift, LNAI 4850* (pp. 53-62). Heidelberg: Springer-Verlag.
19. Vernon, D. (2010). Enaction as a conceptual framework for developmental cognitive robotics. *Journal of Behavioral Robotics*, 1(1), 89-98.
20. Vernon, D., Metta, G., & Sandini, G. (2010). Embodiment in cognitive systems: on the mutual dependence of cognition and robotics. In J. Gray & S. Nefti-Meziani (Eds.), *Embodied Cognitive Systems*. Institution of Engineering and Technology (IET), UK.
21. Gibson, J. J. (1977). The theory of affordances. In R. Shaw & J. Bransford (Eds.), *Perceiving, Acting, and Knowing: Toward an Ecological Psychology* (pp. 67-82). Hillsdale, NJ: Lawrence Erlbaum.
22. Geng, T., Wilson, J., Sheldon, M., Lee, M., & Hulse, M. (2013). Synergy-based affordance learning for robotic grasping. *Robotics and Autonomous Systems*, 61(12), 1626-1630.
23. Pfeifer, R., Verschure, P. (1992). Distributive adaptive control: a paradigm for designing autonomous agents. In F. J. Varela & P. Bourguine (Eds.), *First European Conference on Artificial Life* (pp. 21-30).
24. Verschure, P. F. M. J. (2012). Distributed adaptive control: A theory of the mind, brain, body nexus. *Biologically Inspired Cognitive Architectures*, 1, 55-72.
25. Verschure, P. F. M. J., Voegtlin, T., & Douglas, R. J. (2003). Environmentally mediated synergy between perception and behaviour in mobile robots. *Nature*, 425(6958), 620-624.

26. Demiris, Y., & Meltzoff, A. (2008). The Robot in the Crib: A developmental analysis of imitation skills in infants and robots. *Infant and Child Development*, 17(1), 43-53.
27. Cangelosi, A., & Schlesinger, M. (2015). *Developmental Robotics: from babies to robots*. Cambridge MA, London: The MIT Press.
28. Rizzolatti, G., Fadiga, L., Gallese, V., & Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3, 131-141.
29. Rizzolatti, G., & Craighero, L. (2004). The mirror-neuron system. *Annual Review of Neuroscience*, 27, 169-192.
30. Oztop, E., Kawato, M., & Arbib, M. (2006). Mirror neurons and imitation: a computationally guided review. *Neural Networks*, 19, 254-271.
31. Shanahan, M. (2010). *Embodiment and the Inner Life*. Oxford, New York: Oxford University Press.
32. Csibra, G., & Gergely, G. (2007). Obsessed with goals: Functions and mechanisms of teleological interpretation of actions in humans. *Acta Psychologica*, 124, 60-78.
33. Gilbert, N. (2008). *Agent-Based Models*. Sage Publications, Inc.
34. Richetin, J., Sengupta, A., Perugini, M., Adjali, I., Hurling, R., Greetham, D., & Spence, M. (2010). A micro-level simulation for the prediction of intention and behavior. *Cognitive Systems Research*, 11(2), 181-193.
35. Wolfram, S. (1983). Statistical mechanics of cellular automata. *Reviews of Modern Physics*, 55(3), 601-644.
36. Tisue, S., & Wilensky, U. (2004). NetLogo: Design and Implementation of a Multi-Agent Modeling Environment. In *Agent-2004*.
37. Rao, A., & Georgoff, M. (1995). *BDI Agents: from theory to practice*. Melbourne: Australian Artificial Intelligence Institute.
38. Wooldridge, M. (2000). *Reasoning About Rational Agents*. Cambridge, MA: The MIT Press.
39. Davis, D., & Gwatkin, J. (2010). robo-CAMAL: A BDI motivational robot. *Journal of Behavioral Robotics*, 1(1), 116-129.
40. Lincoln, N., & Veres, S. (2013). Natural language programming of complex robotic BDI agents. *Journal of Intelligent and Robotic Systems*, 71(2), 211-230.
41. TheDmel (Poster). (2010, Sept. 15). Aggressive Quadrotor Part II [Video]. Retrieved from https://youtu.be/geqip_0Vjec
42. InfoLeak (Poster). (2011, Mar. 29). Quadrocopter Tennis [Video]. Retrieved from <https://youtu.be/IVQIglBZo6s>
43. Mazhar, Ali (Poster). (2013, Mar. d25). Skillfull Manipulation [Video]. Retrieved from <https://youtu.be/IWOGtszdSoQ>
44. Robonaut. Retrieved 22 June 2015 from <http://robonaut.jsc.nasa.gov>
45. Hinton, G. E. (2007). Learning multiple layers of representation. *Trends in Cognitive Sciences*, 11(10), 428-434.
46. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S. & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529-533.
47. Pfeifer, R., Scheier, C. (1999). *Understanding Intelligence*. Cambridge, MA: The MIT Press.
48. De Kamps, M. (2012). Towards truly human-level intelligence in artificial applications. *Cognitive Systems Research*, 14(1), 1-9.
49. Mori, M. (1970). Bukimi no tani (the uncanny valley). *Energy*, 7, 33-35.

50. Moore, R. K. (2012). A Bayesian explanation of the Uncanny Valley effect and related psychological phenomena. *Nature Scientific Reports*, 2(864).
51. Maxwell, B. A. (2007). Building robot systems to interact with people in real environments. *Autonomous Robots*, 22, 353-367.
52. Scheutz, M., Schermerhorn, P., Kramer, J., & Anderson, D. (2007). First steps toward natural human-like HRI. *Autonomous Robots*, 22, 411-423.
53. Herrmann, G., & Melhuish, C. (2010). Towards safety in human robot interaction. *International Journal of Social Robotics*, 2(3), 217-219.
54. Moore, R. K. (2015). From talking and listening robots to intelligent communicative machines. In J. Markowitz (Ed.), *Robots That Talk and Listen*. Boston, MA: De Gruyter.
55. Liao, S.-H. (2015). *Awareness and Usage of Speech Technology*. Masters thesis, Dept. Computer Science, University of Sheffield.
56. Introduction and Overview of W3C Speech Interface Framework, <http://www.w3.org/TR/voice-intro/>.
57. Maturana, H. R., Varela, F. J. (1987). *The Tree of Knowledge: The Biological Roots of Human Understanding*. Boston, MA: New Science Library/Shambhala Publications.
58. Cummins, F. (2014). Voice, (inter-)subjectivity, and real time recurrent interaction. *Frontiers in Psychology*, 5, 760.
59. Bickhard, M. H. (2007). Language as an interaction system. *New Ideas in Psychology*, 25(2), 171-187.
60. Cowley, S. J. (Ed.). (2011). *Distributed Language*. John Benjamins Publishing Company.
61. Fusaroli, R., Raczaszek-Leonardi, J., Tylén, K. (2014). Dialog as interpersonal synergy. *New Ideas in Psychology*, 32, 147-157.
62. Scott-Phillips, T. (2015). *Speaking Our Minds: Why human communication is different, and how language evolved to make it special*. Palgrave MacMillan.
63. Powers, W. T., Clark, R. K., & McFarland, R. L. (1960). A general feedback theory of human behavior: part I. *Perceptual and Motor Skills*, 11, 71-88.
64. Powers, W. T., Clark, R. K., & McFarland, R. L. (1960). A general feedback theory of human behavior: part II. *Perception and Motor Skills*, 11, 309-323.
65. Powers, W. T. (1973). *Behavior: The Control of Perception*. NY: Aldine: Hawthorne.
66. Bourbon, W. T., & Powers, W. T. (1999). Models and their worlds. *Int. J. Human-Computer Studies*, 50(6), 445-461.
67. Mansell, W. (2011). Control of perception should be operationalised as a fundamental property of the nervous system. *Topics in Cognitive Science*, 3, 257-261.
68. Powers, W. T. (2008). *Living Control Systems III: The Fact of Control*. Benchmark Publications.
69. DiStefano III, J. J., Stubberud, A. R., & Williams, I. J. (1990). *Feedback and Control Systems*. Schaums Outline Series (2nd ed.). New York: McGraw-Hill.
70. Dalchow, L. (2012). *Subsumption vs Perceptual-Control Architectures in Behaviour-Based Robotics*. Dissertation, University of Sheffield.
71. Taylor, M. M. (1999). Perceptual Control Theory and its application. *Int. J. Human-Computer Studies*, 50, 433-444.
72. Powers, W. T. (1979). The nature of robots - part 1: defining behavior. *Byte*, 4(6).
73. Powers, W. T. (1979). The nature of robots - part 2: simulated control systems. *Byte*, 4(7).
74. Powers, W. T. (1979). The nature of robots - part 3: a closer look at human behavior. *Byte*, 4(8).

75. Powers, W. T. (1979). The nature of robots - part 4: looking for controlled variables. *Byte*, 4(9).
76. Kennaway, J. R. (1999). Control of a multi-legged robot based on hierarchical perceptual control theory. *J. Perceptual Control Theory*, 1(1).
77. Kennaway, J. R. (2004). A Simple and Robust Hierarchical Control System for a Walking Robot. *Unpublished*. Retrieved from <http://www2.cmp.uea.ac.uk/~jrk/Robotics/rk-c2004.pdf>
78. <http://www.perceptualrobots.com>
79. Moore, R. K. (2007). Spoken language processing: piecing together the puzzle. *Speech Communication*, 49(5), 418-435.
80. Lombard, E. (1911). Le sign de l'elevation de la voix. *Ann. Maladies Oreille, Larynx, Nez, Pharynx*, 37, 101-119.
81. Lindblom, B. (1990). Explaining phonetic variation: a sketch of the H&H theory. In W. J. Hardcastle & A. Marchal (Eds.), *Speech Production and Speech Modelling* (pp. 403-439). Kluwer Academic Publishers.
82. Pieraccini, R. (2012). *The Voice in the Machine*. MIT Press, Cambridge, MA.
83. Moore, R. K. (2007). PRESENCE: A human-inspired architecture for speech-based human-machine interaction. *IEEE Trans. Computers*, 56(9), 1176-1188.
84. Moore, R. K., & Nicolao, M. (2011). Reactive speech synthesis: actively managing phonetic contrast along an H&H continuum. *17th International Congress of Phonetics Sciences (ICPhS)*. Hong Kong.
85. Nicolao, M., Latorre, J., & Moore, R. K. (2012). C2H: A computational model of H&H-based phonetic contrast in synthetic speech. *INTERSPEECH*. Portland, USA.
86. Moore, R. K. (2010). Cognitive approaches to spoken language technology. In F. Chen & K. Jokinen (Eds.), *Speech Technology: Theory and Applications* (pp. 89-103). New York Dordrecht Heidelberg London: Springer.
87. Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Phil. Trans. R. Soc. B*, 364(1521), 1211-1221.
88. Weiss, I., & Ray, M. (2001). Model-based recognition of 3D objects from single images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(2), 116-128.
89. Radfar, M. H., Dansereau, R. M., & Sayadiyan, A. (2008). Speaker-independent model-based single channel speech separation. *Neurocomputing*, 72(1-3), 71-78.
90. Gales, M., & Young, S. J. (2007). The application of hidden Markov models in speech recognition. *Foundations and Trends in Signal Processing*, 1(3), 195-304.
91. Moore, R. K. (2014). Interpreting intentional behaviour. In M. Muller, S. S. Narayanan, & B. Schuller (Eds.), *Dagstuhl Seminar 13451 on Computational Audio Analysis* (Vol. 3). Dagstuhl, Germany.
92. Do, C. B., Batzoglu, S. (2008). What is the expectation maximization algorithm? *Nature Biotechnology*, 26, 897-899.
93. Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. Cambridge, MA: The MIT Press.
94. Barto, A. G., Sutton, R. S., & Anderson, C. W. (1983). Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-13(5), 834-846.
95. Grush, R. (2004). The emulation theory of representation: motor control, imagery, and perception. *Behavioral and Brain Sciences*, 27, 377-442.
96. Wilson, M., & Knoblich, G. (2005). The case for motor involvement in perceiving conspecifics. *Psychological Bulletin*, 131(3), 460-473.

97. Pickering, M. J., & Garrod, S. (2013). Forward models and their implications for production, comprehension, and dialogue. *Behavioral and Brain Sciences*, 36(04), 377-392.
98. Chin, R. T., & Dyer, C. R. (1986). Model-based recognition in robot vision. *ACM Computing Surveys*, 18(1), 67-108.
99. Demiris, Y., Khadhour, B. (2005). Hierarchical, Attentive, Multiple Models for Execution and Recognition (HAMMER). *IEEE International Conference on Robotics and Automation (ICRA)*. Barcelona, Spain.
100. Liberman, A., Cooper, F. S., Shankweiler, D. P., & Studdert-Kennedy, M. (1967). Perception of the speech code. *Psychological Review*, 74, 431-461.
101. Skipper, J. I. (2014). Echoes of the spoken past: how auditory cortex hears context during speech perception. *Phil Trans R Soc B*, 369(1651).
102. Kuhl, P. K., Ramirez, R. R., Bosseler, A., Lin, J.-F. L., & Imada, T. (2014). Infants' brain responses to speech suggest Analysis by Synthesis. *Proceedings of the National Academy of Sciences*.
103. Moore, R. K. (2014). Spoken language Processing: Time to Look Outside? In *2nd International Conference on Statistical Language and Speech Processing (SLSP 2014)*. Grenoble.
104. Iacoboni, M. (2005). Understanding others: imitation, language, empathy. In S. Hurley & N. Chater (Eds.), *Perspectives on Imitation: From Mirror Neurons to Memes* (Vol. 1, pp. 255-282). MIT Press.
105. Premack, D., & Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, 1(04), 515-526.
106. Marken, R. S. (2013). Making inferences about intention: perceptual control theory as a 'theory of mind' for psychologists. *Psychological Reports: Measures & Statistics*, 113(1), 257-274.
107. Maslow, A. H. (1943). A theory of human motivation. *Psychological Review*, 50, 370-396.
108. Oudeyer, P.-Y., & Kaplan, F. (2007). What is Intrinsic Motivation? A Typology of Computational Approaches. *Frontiers in Neurorobotics*, 1, 6.
109. Scherer, K. R., Schorr, A., & Johnstone, T. (2001). *Appraisal Processes in Emotion: Theory, Methods, Research*. New York and Oxford: Oxford University Press.
110. Marsella, S., Gratch, J., & Petta, P. (2010). Computational models of emotion. *A Blueprint for Affective Computing-A Sourcebook and Manual*, 21-46.
111. Mehrabian, A. (1996). Pleasure-arousal-dominance: A general framework for describing and measuring individual differences in temperament. *Current Psychology: Developmental, Learning, Personality, Social*, 14, 261-292.
112. Carver, C. S., & Scheier, M. F. (1998). *On the Self-Regulation of Behavior*. Cambridge: Cambridge University Press.