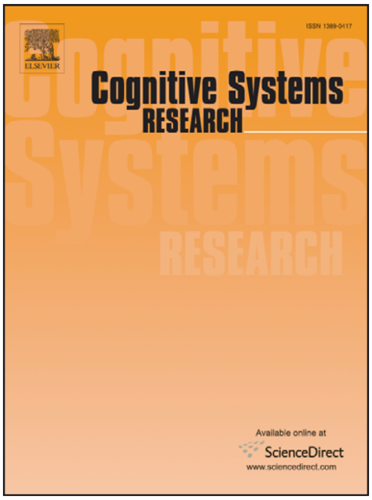

Steps to a "Properly Embodied" Cognitive Science

Mog Stapleton

PII: S1389-0417(12)00024-1
DOI: http://dx.doi.org/10.1016/j.cogsys.2012.05.001
Reference: COGSYS 397

To appear in: *Cognitive Systems Research*

Received Date: 16 April 2011
Accepted Date: 11 May 2012

Please cite this article as: Stapleton, M., Steps to a "Properly Embodied" Cognitive Science, *Cognitive Systems Research* (2012), doi: http://dx.doi.org/10.1016/j.cogsys.2012.05.001

Steps to a "Properly Embodied" Cognitive Science

Mog Stapleton, School of Philosophy, Psychology and Language Sciences, University of
Edinburgh, Dugald Stewart Building, 3 Charles Street, George Square, Edinburgh EH8 9AD,
UK

Abstract

Cognitive systems research has predominantly been guided by the historical distinction between emotion and cognition, and has focused its efforts on modelling the "cognitive" aspects of behaviour. While this initially meant modelling only the control system of cognitive creatures, with the advent of "embodied" cognitive science this expanded to also modelling the interactions between the control system and the external environment. What did not seem to change with this embodiment revolution, however, was the attitude towards affect and emotion in cognitive science. This paper argues that cognitive systems research is now beginning to integrate these aspects of natural cognitive systems into cognitive science proper, not in virtue of traditional "embodied cognitive science", which focuses predominantly on the body's gross morphology, but rather in virtue of research into the interoceptive, organismic basis of natural cognitive systems.

Keywords**:** Proper embodiment; Embodied cognition; Affective neuroscience; Interoception; Internal robotics; Enactive cognitive science.

Corresponding author.
Email address: M.L.Stapleton@sms.ed.ac.uk (M. Stapleton).

1. Introduction

The cognitive science of the twentieth century, reflecting the focus of the individual cognitive sciences, was predominantly interested in perception, memory, problem solving, planning, and other "cognitive" activities. For the most part, researchers interested in cognition ignored those aspects that, as a hangover from Cartesian dualism, were considered "subjective", such as consciousness and affect. This was due to a variety of factors, including the inheritance of behaviourist and cognitivist psychology. Although cognitivism was a reaction to the behaviourism of the early to mid twentieth century and thus directly opposed to many of its claims, they both shared the assumption that the emotional domain was separate from the cognitive domain, and furthermore, that emotion was potentially dissociable from cognition. As a result, cognitive scientists have tended to consider it unnecessary to understand affect in order to understand the other aspects of cognition, and for the most part have left research in this area to a handful of "affective" neuroscientists and psychologists.

My aim in this paper is twofold. Firstly, to give a brief overview of some of the current research in the individual cognitive sciences that suggests that the relation between affect and cognition is more complex, and more important, than has traditionally been held to be the case by mainstream cognitive science. I focus on research from neuroscience and robotics, as I think that these provide the clearest models of current understanding. My second aim in the paper is to show that this work is pointing us in the direction of a new type of embodied cognitive science. Traditional "embodied" cognitive science, whose main focus was on gross morphological sensorimotor interaction ignored the interactions between the control system and the internal body, and thus had no place for the role of affect. I argue that recent work in neuroscience and robotics suggests that cognitive systems are not merely superficially embodied in the sense that the sensorimotor interactions with the environment are the only interactions relevant to cognitive behaviour, but that cognitive systems are “properly embodied”; the internal body matters to cognition.

2. Beyond morphological embodiment

What started in the late twentieth century as an embodied cognitive science revolution, has slowly been becoming mainstream cognitive science (see, for example, Gibbs 2005; Chemero 2009; Shapiro 2011; Barrett 2011). Given the focus on the body, however, it might seem surprising that until recently there was little research involving the internal and affective body beyond designing robots to detect and display emotional expressions so as to make them more appealing to, and facilitate interaction with humans (see, for example, Kismet n.d.). The focus of embodied cognitive science is on decentralizing cognition and modelling how the morphology of the body and its activity often reduce the processing load of the brain (Clark 1997, 2008; Pfeiffer 2007). The paradigmatic examples of this approach to cognitive systems research are Rodney Brooks’s “animats” (Brooks 1991; see also Meyer & Wilson 1991) and Barbara Webb’s robotic crickets (1994, 1996). Brooks’s animats are distinctive in having no central controller as such; rather, their cognitive architecture is organized in layers, each relating to the control of parts of the robot’s body, and which feed back to one another. Similarly, Webb’s robotic crickets model the use of the bodily architecture in cricket phonotaxis, which allows direct response to mating calls, avoiding the need for complex information processing. Good introductions to embodied cognitive science can be found in Clark (1997, 2001, 2003, 2008) and Bermúdez (2010, sect. 4).

The principal way in which bodily morphology engenders and shapes cognitive activity and processes is through sensorimotor interaction with the world. In this way Webb’s robotic crickets exploit their morphology to allow sensing to directly control motor processes. The basic idea is that some of the computational work essential to cognition can be partially offloaded to, and realized by, bodily processes and structures external to the central nervous system. (Physical gestures are another oft-cited example: see Clark 2008; Goldin-Meadow 2005.) Cognition is thus “extended” so that it encompasses parts of the body (and plausibly also those parts of the non-biological world) that support the appropriate offloading of computations. However, this means that as far as standard embodied cognitive science is concerned, the body qua body does not play a special role; only the body in virtue of its ability to be a vehicle of computations. The result is that, although research in this paradigm is based on the role of the body in cognition, the body really isn’t the important factor.

Recently, sensorimotor cognitive science has begun to bring research from developmental psychology together with robotics, and in the process has taken embodiment research towards using not only exteroceptive sensory information (vision, hearing, touch, etc.) to guide action but also proprioceptive and kinaesthetic information (the sense of the location of the body and the movement of the body, respectively). The result is that these sources of internal sensory information not only aid sensorimotor activity but can also be integrated into “higher” cognitive activity. This can be seen specifically with research using the iCub in the European project “Robotcub” and in the follow-up projects such as ITALK and AMARSi. The iCub is designed, with the help of some learning algorithms, to develop roughly like a toddler: through active engagement with its environment, including interaction with humans. An example of what the iCub can do through ontogenetic learning and development can be seen in Morse et al.'s work with the iCub (2010), wherein it has learnt to name objects by associating the name of the object not only with the object of attention when the name is given, but also with the part of egocentric space that the object is normally presented in (for a video of this in action see Barras 2010).

These steps towards a robotics which integrates research from developmental psychology as well as neuroscience are being interwoven with a train of cognitive systems research whose roots lie in the biological sciences and phenomenology rather than the computationalist/functionalist tradition that was the main voice of twentieth century cognitive science. While sensorimotor research in robotics and philosophy of cognitive science has often come to be labelled as "enactive" (principally through Alva Noë's use of the term "enactivism" to describe his sensorimotor theory of consciousness: see Noë 2004, 2009), there is more to enactivism than sensorimotor skills, and as such, enactive cognitive science should not be conflated with sensorimotor cognitive science (for extended arguments on this see Ziemke 2007, 2008; Morse et al. 2011; Di Paolo 2009; Ward & Stapleton forthcoming).

The main distinguishing feature of enactive cognitive science is the focus, not only on the interaction between cognitive systems and their environment (i.e. predominantly sensorimotor interactions), but also on the constitution of cognitive systems and the relation between their constitution and their interaction with the environment (this distinction is from Moreno et al. 2008, cited in Ziemke 2008). For example, Ziemke (2008), and Ziemke and Lowe (2009) have argued that being physical systems which can interact with their environment through sensors and actuators is not sufficient for cognitive embodiment. They propose looking to the bodies of living organisms for the future direction of cognitive robotics: what Di Paolo (2003) refers to as an “organismically inspired robotics”. While enactivism originated in the work on autopoiesis and the structural coupling between an autopoietic system and the environment (see Maturana & Varela 1992; Varela, Thomson & Rosch 1991; Thompson 2007), the principal notions which have come to be of central importance to current cognitive science research are autonomy and adaptivity (Di Paolo 2005; Thompson 2007). Autonomy and adaptivity take the key insights from cellular autopoiesis, such as operational closure, self-construction, and sense-making, and abstract away from the biological implementation. Having said this, the abstraction is not so great that key biological functions such as homeostasis are ignored; the internal is key to enactivism (see Di Paolo 2010).

These baby steps towards an enactive cognitive science can be seen in Vernon’s conceptual framework for the iCub architecture (Vernon 2010), in which the role of cognition is taken to be “to anticipate events and increase the space of actions in which a system can engage” (Vernon 2010, p. 91). More specifically, the position is that:

> (a) cognition is the process by which an autonomous self-governing agent acts effectively in the world in which it is embedded, that (b) the dual purpose of cognition is to increase the agent’s repertoire of effective actions and its power to anticipate the need for and outcome of future actions, and that (c) development plays an essential role in the realization of these cognitive capabilities. (Vernon 2010, p. 90)

The explicit focus on anticipation in this framework is very much in line with the current understanding of the role of prediction in neuroscience, and indeed Vernon claims that:

> […] cognition arises from an agent’s need to compensate for latencies in neural processing by anticipating what may be about to happen and by preparing its actions accordingly. So we can agree fairly easily what cognition is — a process of anticipating events and acting appropriately and effectively — and why it is necessary — to overcome the physical limitations of biological brains and the limitations of bodily movements operating in a dynamic environment. (Vernon 2010, p. 90)

In Vernon’s model the internal components work together to comprise different “cognitive” systems, such that the perception system comprises exogenous salience, endogenous salience, egosphere, and attention selection; the action system comprises gaze control, vergence, reach and grasp, locomotion; anticipation and adaptation are underpinned by the episodic and procedural memory components; motivations are underpinned by the affective state component, which works with the action selection component and provides “a very simple homeostatic process which regulates the autonomous behaviour of the iCub” (Vernon 2010, p. 95). While Vernon’s architecture may be more modular that many embodied and enactive cognitive scientists would be happy with, we can see a progression from more standard architectures in that internal and affective information have critical roles. We see affect as a part of dynamic feedback re-entrant couplings (rather than a feed-forward network) such that “affective information” is feeding directly into action selection and from there into procedural memory, and from there to both gaze control and episodic memory. Even though it might look at first glance as though there is an affect “module”, in fact the cognitive behaviour is a result of the dynamic behaviour between the components, and even the explicitly affective information (i.e. that which the affective component processes and integrates) is feeding back and through many of the components whose principal activity is underpinning perception, action, and anticipation.

3. Affective perception

What evidence do we have that affective information feeds into the kinds of processes Vernon outlines in his cognitive architecture? One source of evidence is from a model of affective predictions in object perception by Barrett and Bar (2009). Barrett and Bar put together research on visual processing in light of the generalized predictive coding approach to neuroscience. At the heart of the generalized predictive coding approach is the hypothesis that the brain is essentially a prediction engine, and the information that we garner from the world is encoded in the errors in these predictions. The brain continues to recalibrate and generate new predictions until the incoming sensory states match those predictions (Bar 2009; Friston 2009; Friston & Kiebel 2009). Prediction has recently become something of a unifying framework guiding understanding at various levels in neuroscience, from the statistics of neural firing to the level of us as agents interacting in the world; see Clark (forthcoming) for an accessible introduction to these principles.

Barrett and Bar address prediction somewhere in between these levels. Their thesis is that object perception is generated by — and through (I add this because the processing is importantly not strictly sequential, but involves a lot of feeding back at various stages) — predictions about the relevance of an object or class of object, that is, its value to the agent either generally or at this particular moment in time. This means that rather than perception being a matter of “bottom-up” processing where the details are put together stage by stage to make the whole, the overall prediction, i.e. the gist of the situation, is processed early on, becoming more and more detailed or accurate through the recurrences. Barrett and Bar use an illustrative analogy of the Dutch style of painting in the sixteenth and seventeenth centuries: first the gist of a situation is sketched, then over time, through the recursive application of ever smaller dabs of paint, a detailed picture emerges. The recursive (and ever finer) dabs of paint in this example correspond to the recursive predictions that are generated as a result of errors in the predictions of sensory states. Their thesis is that object perception arises partly as a result of predictions about the value of that object to the agent.

Drawing on research from Aude Oliva’s computational visual cognition lab at MIT (see, for example, Oliva & Torralba 2006), Barrett and Bar propose that the brain quickly makes an initial prediction about an object using low spatial frequency visual information, and then the details are filled in by memory guided by context. Direct projections between the visual cortex and areas of the prefrontal cortex provide a pathway for this recursive (re-)creation of the visual experience of the object. The previous knowledge which is used to flesh out the gist of the prediction is encoded in sensorimotor patterns which are stored for future use. Importantly for us here, they argue that sensorimotor patterns are sensory in the fullest sense of the term: they not only involve external sensations and their relations to actions, but also internal sensations — from organs, muscles and joints, and how external sensations have influenced these internal sensations (Barrett & Bar 2009, p. 1325). They thus show that the connections between various brain areas give us reason to believe that representations of internal bodily (autonomic and endocrine) changes are part of visual processing right from the stage at which the gist of a situation is being processed by the frontal systems, giving even perception at this paucity of specificity an affective flavour which helps code the relevance/value of the object of perception.

Looking at their model of visual processing in a bit more detail we can see exactly how they propose that affective information feeds into object perception. Visual information comes through the lateral geniculate nucleus (part of the thalamus), at which point a very unspecific "gist" of this information is sent through the fast magnocellular pathway through the dorsal visual stream, which includes the lateral parietal cortex, and also through fast magnocellular pathways to V1 and from there to the dorsal stream. The dorsal stream sends information on to the medial orbitofrontal cortex (mOFC) which then sends information to (i) the autonomic and endocrine systems to effect bodily changes including preparation for action, and (ii) information about those changes that have been ordered to the lateral parietal cortex, feeding that information back into the dorsal stream. This shows that the processing of gist information is affective as the internal bodily changes are caused and the representations of these are fed back into the lateral OFC helping to refine the gist each time with the information about affective value that these carry. The idea is that each time round the processing loops, better and better predictions are being made and the perception of the object is getting less and less gist-like — and at the same time developing more and more meaning (in terms of biological relevance) for the agent in virtue of the affective aspect of the perception. Highly specific visual information (as opposed to "gist" information) gets sent on a different route towards the orbital frontal cortex. From the lateral geniculate nucleus it gets sent through slower parvocellular pathways to both the ventral visual stream and V1, and from there to the ventral stream. Information from the ventral stream gets sent to the lateral OFC (rather than the medial OFC as was the case in the dorsal loop). Also feeding into the lateral OFC is information from the external senses and from the internal bodily changes that were effected as a result of processing in the medial OFC. The lateral OFC thus serves as an association area of all of this information from various senses including interoception. So even the more specific visual processing that builds upon the gist that is being created as a result of the dorsal loops is laden with affective value. I suggest that Barrett and Bar’s model of visual processing looks something like this:

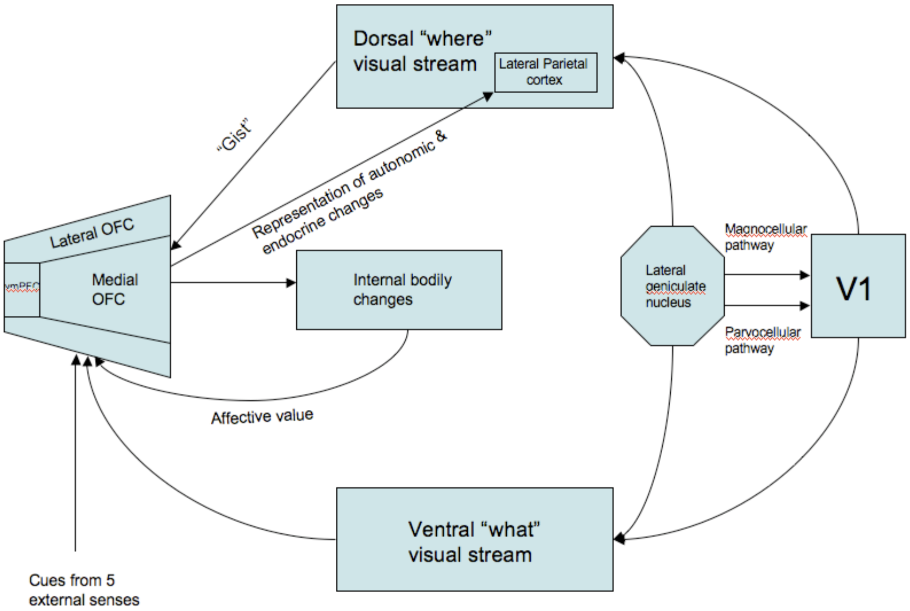

Fig 1: Pictorial representation of Barrett and Bar’s model of affective predictions in object perception (2009)

As my diagram above suggests, upon Barrett and Bar’s proposal all the processing described here is very recurrent and not at all static or sequential in a strict sense. Even in the simplified form that I have presented Barrett and Bar’s proposal, it is clear that affective value is feeding in at various levels (i.e. both dorsal and ventral) and becoming part of the very dorsal processing that it is feeding into. They suggest that it is likely that the real picture is even more complex, which would only further the argument that affective value is inherently a part of visual processing:

> Taken together, these findings indicate that it may be more appropriate to describe the affective predictions generated by the medial and lateral OFC as phases in a single affective prediction evolving over time,

> rather than as two separate 'types' of affective predictions (with one informing the other). This interpretation is supported by the observation that the medial and lateral OFC are strongly connected by intermediate areas; in addition, the lateral OFC receives some low spatial frequency visual information and the medial OFC some high spatial frequency information; and, magnocellular and parvocellular projections are not as strongly anatomically segregated as was first believed (for a review, see Laylock et al. (2007)). Furthermore, there are strong connections throughout the dorsal 'where' and ventral 'what' streams at all levels of processing (Merigan and Maunsell 1993; Chen et al. 2007). Finally, the OFC has widespread connections to a variety of thalamic nuclei that receive highly processed visual input and therefore can not be treated as solely bottom-up structures in visual processing. (Barrett & Bar 2009, p. 1331)

4. The value of the internal

At this point we should step back and consider what is meant by affective. Barrett and Bar use the term to denote information pertaining to the viscera, information which is likely to be a guide to how the system is faring in the world. It is this information which is inherently valenced and value-laden. Typically we think of "valence" as indicating positive or negative experience, but its connection to value can be seen even if we set aside questions of phenomenality. This is clear when we consider the variety of terms for which "valence" is used to refer: hedonic tone; utility; good/bad mood; pleasure/pain; approach/avoidance; rewarding/punishing; appetitive/aversive; and positive/negative (Barrett 2006, p. 40). Valence is rooted in the concept of value (see Colombetti 2005 for a detailed discussion of value and valence), and Barrett and Bliss-Moreau (2009) gesture towards this when — appealing to research by Owren and Rendall (1997, 2001) — they suggest that core affect represents a basic kind of psychological meaning:

> The basic acoustical properties of animal calls (and human voices) directly act on the nervous system of the perceiving animal to change its affective state and in so doing conveys the meaning of the sound. (Barrett & Bliss-Moreau 2009, p. 172)

While I don't want to pursue the issue of whether meaning can be reduced to affective changes, we could instead think of meaning here as being the appraisal of value to the organism. In emotion theory "appraisal" has traditionally had cognitive connotations (see Scherer 1999 for a comprehensive review of appraisal theory) but contemporary appraisal theory acknowledges appraisals which are very low-level and not grounded in deliberation. Scherer writes in regard to low-level appraisals, which make up one of the components of his component theory of emotion:

> … one can argue that we need a general, overarching term to cover the fundamental fact that it is not the objective nature of a stimulus but the organism's "evaluation" of it that determines the nature of the ensuing emotion. A completely automatic, reflexive defence reaction of the organism also constitutes an intrinsic assessment, a valuation, of the noxiousness of the stimulus (although it may not necessarily produce a fully fledged emotion […]). Even if simple feature detection is involved the outcome of the process constitutes an assessment of the significance of the detected stimulus to the organism, given that feature detectors that have any behavioral consequences are automatically "significance detectors". (Scherer 1999, p. 647)

In this context, therefore, appraising value to the system can be something as basic as a response required for maintaining homeostasis (the internal balance that keeps a system viable). In this minimal sense, homeostatic behaviours such as withdrawing from a painful stimulus or seeking water when thirsty are results of an appraisal that the current situation is incompatible with homeostatic viability. There is a temptation to think of the behaviour as being *a result of* the the interoceptive information, and thus a result of some personal or subpersonal level "cognitive" deliberation, however recent research by A. D. Craig on "homeostatic emotions" (Craig 2003a, 2003b) suggests that the action/behaviour is an integral part of interoception (the afferent homeostatic pathway). For example, in the basic pain pathway common to primates and non-primates which rises through the brainstem, the limbic motor cortex (ACC) is directly involved in the loop receiving projections from the medial dorsal nucleus of the thalamus and sending projections on to the periaqueductal grey. And, in the primate specific pathway, the limbic motor cortex is also activated in virtue of direct projections from lamina I, and subsequently projects on to the right anterior insula, in addition to area 3a of the sensorimotor cortex (which projects directly to the primary motor cortex) receiving corollary projections from one of the afferent projections from the thalamus to the interoceptive cortex in the insula (for details, see Craig 2003a, 2003b).

It can be seen from this that the “motor areas” of the central nervous system are part of the very homeostatic loop itself, rather than functioning — at this basic level — as a result of deliberation conceived in either personal or subpersonal level terms. They are so entwined with the afferent homeostatic signals which ground interoception that it looks as though interoception is not merely the passive representation of the physiological changes in the body, but has the motor aspects already factored in. In other words, interoception includes motor information. So if interoception includes information about and preparation for homeostatic behaviour, then it is by nature functioning as a basic appraisal machine adapting the system in response to perturbations from the environment. The point to which I want to draw attention is that if this minimal model of appraisal of value is correct then it is appropriate to understand valence as affective motivation: affective because it is constituted by afferent homeostatic information, and motivational because it is also constituted by the activation of motor areas.

5. Internal robotics

In a similar vein to Craig's work in neuroscience on homeostatic behaviours, Parisi, in his 2004 article "Internal Robotics", explains how his lab have evolved behaviours such as sleeping, feeling pain. and feeling hungry and thirsty in their agents, in order for them to cope with particular environmental problems. The effective evolution of these behaviours illustrates the co-evolution of the internal environment and the control system (a neural network in the robots, the nervous system in organisms) and external interactions.

Parisi argues that the behaviour of organisms is a result of two types of interaction: (1) between the nervous system and the external environment, and (2) between the nervous system and the internal environment. The nervous system can be thought of as a physical system which realizes the function of mapping inputs onto outputs "in ways that allow an organism to survive and reproduce in its environment" (Parisi 2004, p. 326). Inputs are caused by either physical or chemical events outside the nervous system, and likewise outputs cause physical or chemical events outside the nervous system. But (and this is what "embodied cognitive science" tends to overlook) "outside the nervous system" includes not only the external environment but also the internal body. The internal body provides inputs to the nervous system physically through somatosensory and proprioceptive receptors (relating information about movement and location of body) and through chemical means with molecules from the endocrine system modulating the nervous system and even molecules from the nervous system itself feeding back to it. Likewise, the internal body receives outputs from the nervous system both by chemical means and by neural connections to parts of the autonomic nervous system. Of course, internal interactions are predominantly chemical (as opposed to the predominantly physical interactions between the nervous system and the external environment), and chemical interactions have quite different properties to physical interactions. Physical interactions are mediated by neuron-to-neuron connections where the specific “weights” of the connections seem to be the predominant factor in information transfer such that the neurotransmitters that are used in the interactions play a qualitative (rather than quantitative) role. Interaction between the nervous system and the body relies predominantly on molecular based information transfer and can be slower, diffuse and reliant on quantitative effects to activate thresholds. Nevertheless, apart from during states such as sleep (plausibly), the nervous system is constantly interacting with both the external environment and the internal environment (i.e. the internal, autonomic and homeostatic body) as depicted in Fig.2 below.

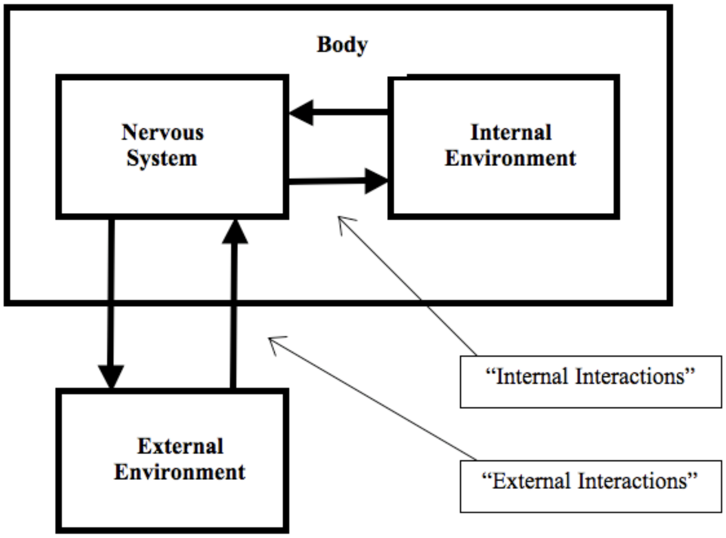

Fig 2: Internal and external interactions (adapted from Parisi 2004)

It might seem appropriate to think of internal interactions (i.e. between the nervous system and the internal environment) as underpinning the emergence of affect, while the external interactions (between the nervous system and the external environment) underpin cognition.

However, Parisi argues that one cannot truly separate the cognitive and affective components, nor understand one without the other:

> For example, an organism may be able to do many different things (cognitive component), but what the organism actually does at any particular time depends on its motivational state (affective component). The current motivational state of an organism is a result of the interactions of the organism's nervous system with the inside of the body, and it controls aspects of behaviour that seem to be purely cognitive, such as selective attention. (Parisi 2004, p. 332)

The moral of Parisi's paper is that the behaviour of organisms results from both internal interactions and external interactions; organisms live in two worlds, the external and the internal. The building of artificial cognitive systems, therefore, whether their purpose is to model natural cognitive systems or whether they are being designed as artificial cognitive systems in their own right, must also simulate both interactions. In his own words:

> The behaviour of an organism results from both the interactions of the organism's nervous system with the external environment and its interactions with the internal environment. Therefore, what is needed is not only an external robotics but also an internal robotics. If we want to understand the behaviour of organisms what we need to reproduce in artificial physical organisms, i.e. robots, is not only the external morphology of an organism's body and the interactions of the organism's nervous system with the external environment, but also the internal physical structure of the organism's body and the interactions of the organism's nervous system with what lies inside the body. (Parisi 2004, p. 326)

6. Affective cognition

The work I have reviewed so far gestures towards a very different way of thinking about enhancing cognitive systems with affectivity than by adding some sort of an emotion module.

Crudely put the argument is that internal bodily affect is crucial to cognitive systems and that an "emotion chip" just won't suffice. This gives us reason to step away from the emotion–cognition distinction as it has traditionally been elucidated in the cognitive sciences. Without denying that there are behaviours and experiences that are usefully labelled using emotion terms and that these episodes can have effects on functions traditionally labelled "cognitive" such as perception, memory, and learning, we can observe that there is another, more fundamental, aspect of affect which is integrated into basic cognitive activity. Recent work by prominent figures in neuroscience, such as Luiz Pessoa, Antonio Damasio and Joseph LeDoux underpins this change of focus.

6.1 Somatic markers don't suffice

Outwith the neuroscience community, Antonio Damasio is best known for his work on somatic markers in which he argues that emotion plays an important role in some cognitive processes (see, for example, Damasio 1994). The somatic marker hypothesis for example is the hypothesis that thoughts that arise get tagged with affective information in the ventromedial prefrontal cortex (vmPFC) which enables the healthy person to utilize information about previous experience quickly and subconsciously during decision making. In the Iowa Gambling Task that is the principal methodology in Damasio and his colleagues' studies, this means that healthy participants learn quickly which decks of cards yield high punishments as well as high pay-offs, and so they naturally gravitate towards the safe decks, which yield lower pay-offs but also lower punishments, and are ultimately better in terms of awarding winnings. The somatic marker hypothesis is based on Damasio and his colleagues' findings that in participants who have damage to the ventromedial prefrontal cortex there is no gravitation towards the safe decks. Measurements of the skin conductance response of both healthy and vmPFC participants showed that vmPFC participants did not have the kind of skin response in anticipation of choosing the risky decks that healthy participants exhibited. Damasio and his colleagues proposed that this was due to the vmPFC being the site where “cognitive” information coming from other cortical areas is “tagged” by “emotional” information coming from limbic areas, and so the disruption of these links caused by vmPFC damage dissociates the affective information from the cognitive information. The hypothesis fits with what we know about the behaviour of those with vmPFC damage: people with this kind of damage often perform well on standard psychological tests, IQ tests, and so forth, and yet appear to be ultimately impaired in normal life; they have difficulty making simple decisions when there is no clear “right” answer such as when choosing what clothes to wear or what restaurant to choose (Saver & Damasio 1991); they also tend to be impulsive and engage in risky behaviours that would have been alien to their pre-damaged self (ibid.).

This hypothesis gives emotion an important place in cognitive processing; however, it still takes emotion to be separate from cognition. It is thoughts, cognitions, that are tagged with affect in virtue of links to what has traditionally been considered the emotion system, i.e. the amygdala, and the insula, which is the locus of representations of the body. In effect, it is not really that emotion plays a part in cognitive processes, but rather that for Damasio these somatic markers solve a kind of localized frame problem: if choosing from one deck of cards in the gambling task has resulted in a negative bodily reaction, negative affect gets tagged to thoughts of that deck, and so when thoughts of that deck arise again for whatever reason that affect is played back, shutting off — or severely reducing — possibilities for action. But making some possibilities for action less motivating (or more motivating, as the case may be) doesn’t mean that emotion is actually playing a part in any of the cognitive processes per se, but that it is playing a role in the meta-cognitive process of using cognitive processes plus affective processes for action selection. So while decision making might be thought of as a cognitive process because it encompasses cognitive processes such as memory, planning, etc., it is not itself a cognitive process in the same manner as those processes. Rather it is a conglomerate of processes, and therefore although it is interesting that it might also require encompassing processes traditionally considered to be affective, this does not need to be a radical claim about cognitive processes per se.

6.2 Integrating affective information

Whilst the somatic marker hypothesis may not support the strong claim that cognition must be affectively embodied, Damasio's more recent work does, in precisely the interoceptive way explicated in the previous sections. Damasio (2010) talks of "primordial feelings" being grounded in brainstem structures, in particular the nucleus tractus solitarius and parabrachial nucleus. These structures are richly connected and are the locus of visceral sensory maps which are key to regulating homeostatic processes. In combination with their connections with the periaqueductal gray and the superior colliculus, Damasio hypothesizes that these brainstem structures instantiate a primordial feeling state in virtue of their being the locus of initial neural representations of changes in the autonomic and endocrine systems. These primordial feelings are constant and provide a background to all cognition such that exteroceptive and interoceptive perception, while having different foci, are nevertheless consistently meshed and thus each feed into both experience and connections to other cognitive processes. This suggests that it may be fruitful for us to think of interoceptive perception as part and parcel of exteroception as well, at least in so far as exteroceptive information is processed. We can see this in particular in Damasio's explanation of the role that the superior colliculus plays in the perceptual pathway.

The superior colliculus (SC) is a primarily visual structure, receiving information directly from both the retina and the visual cortex . However, in addition to these maps of the visual world, the SC also contains “topographical maps of auditory and somatic information, the latter hailing from the spinal cord as well as the hypothalamus” (Damasio 2010, p. 84). Of particular interest is that these maps may all be integrated:

> The three varieties of maps — visual, auditory, and somatic — are in a spatial register. This means that they are stacked in such a precise way that the information available in one map for, say, vision, corresponds to the information on another map that is related to hearing or body state. There is no other place in the brain where information available from vision, hearing, and multiple aspects of body states is so literally superposed, offering the prospect of efficient integration. The integration is made all the more significant by the fact that its results can gain access to the motor system (via the nearby structures in the periaqueductal gray as well as the cerebral cortex). (Damasio 2010, p. 84)

Thus, according to Damasio, even very early on in the neural pathways we find that affective, perceptual and motor information are inseparably intertwined. Due to the non-linear and recurrent nature of the brain’s pathways, such integrated channels of information feed into what we might have thought of as purely perceptual pathways, and this integrated information constrains what is then available for perception.

6.3 The amygdala as supporting the interdependency of affect and cognition

Pessoa’s review of emotion and cognition (2008) gives us good neuroscientific grounding for what might seem to be a radical thesis: that affect and cognition are mechanistically interdependent (and often integrated). He argues that paradigmatic “cognitive” processes such as memory and attention involve information from “emotion”, whether considered in terms of structure, function, or connectivity. While cognitive processes have traditionally been located in the cortex, and correlates of emotion in subcortical (limbic) and para-limbic areas (those areas that used to be considered the old, “mammalian” brain)[1], Pessoa argues that these distinctions are ungrounded. Traditionally what were considered emotional processes such as motivation, drive, appraisals, bodily changes and arousal were thought (i) not to be involved in cognitive processes such as attention, (ii) to be independent of top-down factors, and (iii) to be context-independent. However, Pessoa explains that even the most paradigmatic of the structures that have been associated with emotion, the amygdala, contravenes all three of these motivations.

In respect of connectivity of the amygdala, it is important to note that the amygdala receives from, and gives out information to, areas other than those traditionally regarded as ‘emotional’ or action-provoking. The brain has a small world topology which means that all brain areas are connected by one or two intermediate areas (Sporns et al. 2004; Sporns & Zwi 2004; Bullmore & Sporns 2009). It might be thought that because the prefrontal areas are among the most distant from sensory periphery they receive the most highly processed and integrated information. Highly processed information is supposed to bring greater flexibility, and support the abstract processing required for “cognition”. However, the amygdala is connectively equally removed from the sensory periphery, and so receives just as highly processed and integrated information as the prefrontal areas. Moreover, the amygdala makes (and receives) widespread projections to the rest of the brain. It is one of the most highly connected regions of the brain. So if viewed from a perspective of connectivity, the amygdala is in the “geometric centre” of the topological map (see Pessoa 2008, Fig. 1), and even though it is a core affective region, it is “at least as well situated to integrate and distribute information as certain PFC territories” (Pessoa 2008, p. 151).

---

[1] See Maclean & Kral 1973 for the origins of this triune brain theory.

There thus seems insufficient justification for asserting that one of the key “emotion” areas, the amygdala, is an “affective” rather than a “cognitive” structure. Pessoa (2010) expands on this and suggests that the amygdala should instead be considered a predictive structure involved in situations where the organism must work out the answers to questions in the environment such as “What is it?” and “What’s to be done?”

Understanding the amygdala as a hub in the brain’s networks we can begin to see that it does indeed play an important role in emotional situations. But this role is not of the kind that we imagine when we are entrenched in the conception of the brain as a one-area-one-function machine. This framework resulted in interpreting data in such a way that the amygdala has generally become thought of as a “rapid-response fear module” (for a detailed rejection of this view, see Sander et al. 2003). Now that we have a different framework in theoretical neuroscience in which to view data within, that of networks (Sporns 2010; Sporns & Zwi 2004) and prediction (Bar 2009; Friston & Kiebel 2009), we can see that there is evidence that the amygdala plays a far more important role then previously supposed. The role it plays is of coding for biological relevance (Sander et al. 2003). This can be understood in terms of the amygdala’s function being to “direct the various sources of attention […] towards a source of sensory stimulation (such as an object) when the predictive value of that stimulation is unknown or uncertain” (Barrett et al. 2007).

The literature which shows that the amygdala is involved in emotion processing does not necessarily imply either that the amygdala is an emotion structure or that its activation pertains to emotional rather than non-emotional stimuli. Rather, conceived of as coding for significance, previous findings can be accounted for, while also accounting for the important role it plays in non-emotional processing (that is, processing that would not traditionally be considered "emotional"). Similarly, Pessoa and Adolphs (2010) argue that the amygdala is not an emotion module, but a core brain circuit with "broad connectivity with the cortex and other subcortical structures" enabling it to play a modulatory role in multiple networks:

> The precise functional importance of the amygdala in these networks remains to be investigated, but it is unlikely that it will map specifically onto emotion. Instead, we think that it corresponds to broader and more abstract dimensions of information processing, including processing of salience, significance, ambiguity, unpredictability and other aspects of 'biological value'. More broadly, we argue that the amygdala has a key role in solving the following problem: how can a limited capacity information processing system that receives a constant stream of diverse inputs selectively process those inputs that are the most relevant to the goals of the animal? (Pessoa & Adolphs 2010, p. 780)

There is much more work that needs to be done here to show how this new conception of the activity of the amygdala fits in with the emerging theoretical neuroscience frameworks. But for our purposes the importance of this re-evaluation of the amygdala's role in processing should be clear. If activity in the amygdala is part of a network which codes for biological significance, and it is a hub projecting to, and receiving projections from, most areas in the brain, including those previously deemed to be "cognitive", it is going to be very difficult to continue to hold the coarse distinction between either affect or emotion and cognition in terms of neural processing.

7. Conclusions: Towards a properly embodied cognitive science

Current neuroscience strongly suggests that processing is neither “affective” nor “cognitive”. Clearly there is still going to be a distinction in that certain networks may predominantly underpin certain activities, but these must be specified in each particular case, for example, the neural processing underpinning perception, or fear, or surprise, and so on. If there is such a thing as “affective” or “cognitive” processing this will only be discovered by understanding in full the correlates of these categories of abilities, activities or behaviours. Once this has been done we may compare the correlates of all those categories which we consider “affective” and compare these with those we consider “cognitive”. If the difference reveals a pattern of processing that is particular to affective categories or cognitive categories, then we may have grounds for considering there to be affective or cognitive domains. However, (i) given the evidence that we have so far on the underpinnings of these categories — in particular, the amount (and importance of) recurrency in the neural processing — this looks unlikely to be the case; and (ii) this will depend on whether theses categories have been accurately distinguished as “cognitive” or “affective”.

Even LeDoux, whose work on the amygdala as a fear centre helped to propagate this coarse distinction, has recently shifted his focus away from “emotion” circuits and towards “survival circuits”, information about which feeds into the cognitive workspace along with information from explicit memory, language, environmental activity, body feedback, and central nervous system arousal (LeDoux 2012). This approach is reminiscent of Brooks’s layered robots discussed in section 2, and yet we can see that survival circuits, body feedback, and CNS arousal support a more adaptive picture of cognition; that which underpins flexible adaptive behaviour in an environment. The recent work on affective perception and the integration of affect into "cognitive" functions reviewed here gives a good indication of the importance of the role of internal information in natural cognitive systems and suggest ways in which these aspects might be implemented in artificial cognitive systems. This would extend the concept of embodiment in cognitive science beyond the current sensorimotor embodiment paradigm towards the organismic, enactive paradigm. Such a properly embodied cognitive science embraces the affective not merely as critical for realistic cognitive systems but as integrated in cognition itself.

Damasio has long argued for the importance of homeostasis and interoceptive information for emotion and cognition (Damasio 1994, 1999) and develops this in a highly accessible way in the central chapters of his recent book *Self Comes to Mind* (2010), and this is beginning to be integrated into research in cognitive modelling and robotics, together with research on the physiology of pain and touch. (While touch is usually considered to be an exteroceptive sense there is also an argument for it being categorized as interoceptive; see Craig 2003a, 2003b, 2008.) But while haptics is now a common facet of robotics, interoception has yet to become an orthodox part of cognitive systems research. Work from research groups such as those led by Tom Ziemke (in particular, the Integrating Cognition, Emotion, and Autonomy (ICEA) project) and Ezequiel Di Paolo, who have been working on developing the insights from physiology and neuroscience outlined in section 2 and modelling these in robotic agents, is changing this, however, and robotic modelling is beginning to integrate processes beyond sensorimotor interaction (see Morse et al. 2011).

Interoception is inherently entwined with affect. While there are disputes in the emotion literature as to how much bodily feelings are involved in emotion (if at all) there is little doubt as to their role in affect more generally; and it is quite plausible that the basis of valence and arousal (what Barrett calls "core-affect") lie in interoception (Craig 2008; Barrett & Bliss-Moreau 2009). The move towards more biologically plausible robotics, a robotics which is not only sensorimotor and superficially autonomous but which is interoceptive and provides a way to be deeply autonomous, is a step towards an affective robotics, and thus a step towards a cognitive science which is not merely embodied in terms of its sensorimotor possibilities but "properly embodied".

Acknowledgements: The author would like to thank the reviewers Tom Ziemke and Giovanna Colombetti for their helpful comments, and David Harris for his help with proof reading.

References

Bar, M. (2009). Predictions: a universal principle in the operation of the human brain. Introduction. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *364*(1521), 1181–1182. doi:10.1098/rstb.2008.0321

Barras, C. (2010, September 22). Robots on TV: AI goes back to baby basics. *New Scientist Online*. Retrieved from http://www.newscientist.com/article/dn19470-robots-on-tv-ai-goes-back-to-baby-basics.html

Barrett, L. F. (2006). Valence is a basic building block of emotional life. *Journal of Research in Personality*, *40*(1), 35–55. doi:10.1016/j.jrp.2005.08.006

Barrett, L. (2011). *Beyond the Brain: How Body and Environment Shape Animal and Human Minds*. Princeton, N.J: Princeton University Press.

Barrett, L. F., & Bar, M. (2009). See it with feeling: affective predictions during object perception. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *364*(1521), 1325–1334. doi:10.1098/rstb.2008.0312

Barrett, L. F., & Bliss-Moreau, E. (2009). Affect as a Psychological Primitive. *Advances in Experimental Social Psychology*, *41*, 167–218. doi:10.1016/S0065-2601(08)00404-8

Barrett, L. F., Ochsner, K. N., & Gross, J. J. (2007). On the automaticity of emotion. In J. Bargh (Ed.), *Social Psychology and the Unconscious: The Automaticity of Higher Mental Processes*. New York: Psychology Press.

Bermúdez, J. L. (2010). *Cognitive Science: An Introduction to the Science of the Mind*. Cambridge: Cambridge University Press.

Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, *47*(1–3), 139–159. doi:10.1016/0004-3702(91)90053-M

Bullmore, E., & Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews. Neuroscience*, *10*(3), 186–198. doi:10.1038/nrn2575

Chemero, A. (2009). *Radical Embodied Cognitive Science*. Cambridge, Mass.: MIT Press.

Chen, C.-M., Lakatos, P., Shah, A. S., Mehta, A. D., Givre, S. J., Javitt, D. C., & Schroeder, C. E. (2007). Functional anatomy and interaction of fast and slow visual pathways in macaque monkeys. *Cerebral Cortex*, *17*(7), 1561–1569. doi:10.1093/cercor/bhl067

Clark, A. (1997). *Being There: Putting Brain, Body and World Together Again*. Cambridge, Mass.: MIT Press.

Clark, A. (2001). *Mindware: An Introduction to the Philosophy of Cognitive Science*. New York: Oxford University Press USA.

Clark, A. (2003). *Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence*. Oxford: Oxford University Press.

Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: Oxford University Press.

Colombetti, G. (2005). Appraising Valence. *Journal of Consciousness Studies*, *12*(8-10), 103–126.

Craig, A. D. (2003a). A new view of pain as a homeostatic emotion. *Trends in Neurosciences*, *26*(6), 303–307.

Craig, A. D. (2003b). Interoception: the sense of the physiological condition of the body. *Current Opinion in Neurobiology*, *13*(4), 500–505.

Craig, A. D. (2008). Interoception and Emotion: A Neuroanatomical Perspective. In M. Lewis, J. M. Haviland-Jones, & L. F. Barrett (Eds.), *Handbook of Emotions* (3rd edn., pp. 272–288). New York: Guilford Press.

Dalgleish, T. (2004). The emotional brain. *Nature Reviews. Neuroscience*, *5*(7), 583–589. doi:10.1038/nrn1432

Damasio, A. (1994). *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: G. P. Putnam.

Damasio, A. (1999). *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*. New York: Harcourt Brace. Retrieved from http://lcweb.loc.gov/catdir/toc/99026357.html

Damasio, A. (2010). *Self Comes to Mind: Constructing the Conscious Brain*. New York: Pantheon Books.

De Jaegher, H., Di Paolo, E. A., & Gallagher, S. (2010). Can social interaction constitute social cognition? *Trends in Cognitive Sciences*, *14*(10), 441–447. doi:10.1016/j.tics.2010.06.009

Di Paolo, E. A. (2003). Organismically-inspired robotics: Homeostatic adaptation and natural teleology beyond the closed sensorimotor loop. In K. Murase & T. Asakura (Eds.), *Dynamic Systems Approach for Embodiment and Sociality: From Ecological Psychology to Robotics*, International series on advanced intelligence. Magill, South Australia: Advanced Knowledge International.

Di Paolo, E. A. (2005). Autopoiesis, adaptivity, teleology, agency. *Phenomenology and the Cognitive Sciences*, *4*(4), 429–452. doi:10.1007/s11097-005-9002-y

Di Paolo, E. A. (2008). Extended Life. *Topoi*, *28*(1), 9–21. doi:10.1007/s11245-008-9042-3

Di Paolo, E. A. (2009). *Shallow and Deep Embodiment*. Retrieved from http://shanghailectures.org/guest-lectures/43-presenter/177-ezequiel-di-paolo

Di Paolo, E. A. (2010). Robotics inspired in the organism. *Intellectica*, *53-54*, 129–162.

Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences, 13*(7), 293–301. doi:10.1016/j.tics.2009.04.005

Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences, 364*(1521), 1211–1221. doi:10.1098/rstb.2008.0300

Gendron, M., & Barrett, L. F. (2009). Reconstructing the Past: A Century of Ideas About Emotion in Psychology. *Emotion Review*, *1*(4), 316–339. doi:10.1177/1754073909338877

Gibbs, R. (2005). *Embodiment and Cognitive Science* (1st ed.). Cambridge: Cambridge University Press.

Goldie, P. (2000). *The Emotions: A Philosophical Exploration*. Oxford: Clarendon Press.

Goldin-Meadow, S. (2005). *Hearing Gesture: How Our Hands Help Us Think*. Cambridge, Mass.: Harvard University Press.

Green, D. (1996). *Cognitive Science: An Introduction*. Oxford: Blackwell.

Griffiths, P. (1997). *What Emotions Really Are: The Problem of Psychological Categories*. Chicago: University of Chicago Press.

Harre, R. (2002). *Cognitive Science: A Philosophical Introduction*. London: Sage.

Kismet (n.d.). http://www.ai.mit.edu/projects/humanoid-robotics-group/kismet/kismet.html.

Laycock, R., Crewther, S. G., & Crewther, D. P. (2007). A role for the “magnocellular advantage” in visual impairments in neurodevelopmental and psychiatric disorders. *Neuroscience and Biobehavioral Reviews*, *31*(3), 363–376. doi:10.1016/j.neubiorev.2006.10.003

LeDoux, J. (1998). *The Emotional Brain: The Mysterious Underpinnings of Emotional Life*. London: Weidenfeld & Nicolson.

LeDoux, J. (2007). The amygdala. *Current Biology: CB*, *17*(20), R868–874. doi:10.1016/j.cub.2007.08.005

LeDoux, J. (2012). Rethinking the emotional brain. *Neuron*, *73*(4), 653–676. doi:10.1016/j.neuron.2012.02.004

Lepore, E., & Pylyshyn, Z. (1999). *What is Cognitive Science?* Oxford: Blackwell.

MacLean, P. D., & Kral, V. A. (1973). *A triune concept of the brain and behaviour*. Toronto: Published for the Ontario Mental Health Foundation by University of Toronto Press.

Maturana, H. R., & Varela, F. J. (1992). *Tree of Knowledge: The Biological Roots of Human Understanding*. (Rev. Edn.). New York: Shambhala.

Merigan, W. H., & Maunsell, J. H. (1993). How parallel are the primate visual pathways? *Annual Review of Neuroscience*, *16*, 369–402. doi:10.1146/annurev.ne.16.030193.002101

Meyer, J., & Wilson, S. (Eds.). (1991). *From Animals to Animats: Proceedings of the First International Conference on SImulation of Adaptive Behavior (SAB’90)*. Cambridge, Mass.: MIT Press.

Morse, A. F, Belpaeme, T., Cangelosi, A., & Smith, L. (2010). Thinking With Your Body: Modelling Spatial Biases in Categorization Using a Real Humanoid Robot. *Proceedings of the Cognitive Science Conference*.

Morse, Anthony F., Herrera, C., Clowes, R., Montebelli, A., & Ziemke, T. (2011). The role of robotic modelling in cognitive science. *New Ideas in Psychology*. In press, corrected proof. doi: 10.1016/j.newideapsych.2011.02.001

Murray, L., & Trevarthen, C. (1985). Emotional regulation of interactions between 2-month-olds and their mothers. In Social Perception in Infants. In T. M. Field & N. A. Fox (Eds.), *Social Perception in Infants* (pp. 177–197). Norwood, N.J.: Ablex.

Noë, A. (2004). *Action in Perception*. Cambridge, Mass.: MIT Press.

Noë, A. (2009). *Out of Our Heads: Why You Are Not Your Brain, and Other Lessons from the Biology of Consciousness*. New York: Hill and Wang.

Oatley, K., Keltner, D., & Jenkins, J. M. (2006). *Understanding Emotions*. Oxford: Blackwell.

Oliva, A., & Torralba, A. (2006). Building the gist of a scene: the role of global image features in recognition. *Progress in Brain Research*, *155*, 23–36. doi:10.1016/S0079-6123(06)55002-2

Owren, M. J., & Rendall, D. (1997). An affect-conditioning model of non-human primate vocal signaling. In D. H. Owings, M. D. Beecher, & N. S. Thompson (Eds.), *Communication*. New York: Springer-Verlag.

Owren, M. J., & Rendall, D. (2001). Sound on the rebound: bringing form and function back to the forefront in understanding nonhuman primate vocal signaling. *Evolutionary Anthropology*, *10*(2), 58–71.

Panksepp, J. (2003). At the interface of the affective, behavioral, and cognitive neurosciences: decoding the emotional feelings of the brain. *Brain and Cognition*, *52*(1), 4–14.

Panksepp, J. (2005). On the embodied neural nature of the core emotional affects. *Journal of Consciousness Studies*, *5*, 158–184.

Parisi, D. (2004). Internal robotics. *Connection Science*, *16*(4), 325.

Pessoa, L. (2008). On the relationship between emotion and cognition. *Nature Reviews. Neuroscience*, *9*(2), 148–158. doi:10.1038/nrn2317

Pessoa, L. (2010). Emotion and cognition and the amygdala: from “What is it?” to “What’s to be done?” *Neuropsychologia*, *48*(12), 3416–3429. doi:10.1016/j.neuropsychologia.2010.06.038

Pessoa, L., & Adolphs, R. (2010). Emotion processing and the amygdala: from a “low road” to “many roads” of evaluating biological significance. *Nature Reviews. Neuroscience*, *11*(11), 773–783. doi:10.1038/nrn2920

Pfeifer, R. (2007). *How the Body Shapes the Way We Think: A New View of Intelligence*. Cambridge, Mass.: MIT Press.

Prinz, J. J. (2004). *Gut Reactions: A Perceptual Theory of Emotion* (illustrated edn.). Oxford: Oxford University Press.

Rumesin, H. M., & Varela, F. J. (1992). *The Tree of Knowledge: The Biological Roots of Human Understanding* (3rd revised edn.). New York: Shambhala.

Sander, D., Grafman, J., & Zalla, T. (2003). The human amygdala: an evolved system for relevance detection. *Reviews in the Neurosciences*, *14*(4), 303–316.

Saver, J. L., & Damasio, A. R. (1991). Preserved access and processing of social knowledge in a patient with acquired sociopathy due to ventromedial frontal damage. *Neuropsychologia*, *29*(12), 1241–1249.

Scherer, K. R. (1999). Appraisal Theory. In T. Dalgleish & M. Power (Eds.), *Handbook of Emotion and Cognition* (pp. 637–663). Oxford: Wiley.

Shapiro, L. A. (2011). *Embodied Cognition*. New problems of philosophy. New York: Routledge.

Sloman, A. (1998). Damasio, Descartes, alarms and meta-management. *Systems, Man and Cybernetics*.

Sporns, O. (2010). *Networks of the Brain*. Cambridge, Mass.: MIT Press.

Sporns, O., Chialvo, D. R., Kaiser, M., & Hilgetag, C. C. (2004). Organization, development and function of complex brain networks. *Trends in Cognitive Sciences*, *8*(9), 418–425. doi:10.1016/j.tics.2004.07.008

Sporns, O., & Zwi, J. D. (2004). The small world of the cerebral cortex. *Neuroinformatics*, *2*(2), 145–162. doi:10.1385/NI:2:2:145

Thompson, E. (2007). *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Cambridge, Mass.: Harvard University Press.

Varela, F. J., Thompson, E., & Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, Mass.: MIT Press.

Vernon, D. (2010a). Enaction as a Conceptual Framework for Developmental Cognitive Robotics. *Paladyn Journal of Behavioral Robotics*, *1*(2), 89–98.

Vernon, D. (2010b, January 29). Cognitive development and the iCub humanoid robot. Presented at the EUCogII Conference, ETH Zurich, Switzerland. Retrieved from https://cast.switch.ch/vod/clips/1r94cee789

Ward, D., & Stapleton, M. (2012). Es are good: Cognition as enacted, embodied, embedded, affective and extended. In F. Paglieri (Ed.), *Consciousness in Interaction: The Role of the Natural and Social Context in Shaping Consciousness* (pp. 89–104). John Benjamins.

Webb, B. (1994). Robotic experiments in cricket phonotaxis. In D. Cliff, P. Husbands, J.-A. Meyer, & S. W. Wilson (Eds.), *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. Cambridge, Mass.: MIT Press.

Webb, B. (1996, December). A Cricket Robot. *Scientific American*, 94–99.

Ziemke, T. (2007). The Embodied Self: Theories, Hunches and Robot Models. *Journal of Consciousness Studies*, *14*(7), 167–179.

Ziemke, T. (2008). On the role of emotion in biological and robotic autonomy. *Bio Systems*, *91*(2), 401–408. doi:10.1016/j.biosystems.2007.05.015

Ziemke, T., & Lowe, R. (2009). On the Role of Emotion in Embodied Cognitive Architectures: From Organisms to Robots. *Cognitive Computation*, *1*(1), 104–117. doi:10.1007/s12559-009-9012-0