

Active Perception in Meaningful Worlds*

Michael Wheeler

School of Cognitive and Computing Sciences,
University of Sussex, Brighton BN1 9QH, U.K.

Telephone:+44 273 678524

Fax:+44 273 671320

E-Mail:michaelw@cogs.susx.ac.uk

Abstract

This paper sketches a conceptual framework for thinking about perception. By

further representational description (this time of the desired actions) to which the action-mechanisms then respond. Brooks refers to these principles as *functional decomposition* or the *sense-model-plan-act* framework [7, 8]. In mainstream A.I., perception became essentially a process of transduction followed by inference.

As a key example of a theory of perception embedded in the traditional sense-model-plan-act framework, consider Marr's account of vision [21]. In broad terms, Marr thought that the function of the visual system is to construct three-dimensional internal representations of objects in the visual field from informational clues available in the two-dimensional retinal image. The staged process moves via an intermediate viewer-centred representation of distance and orientation to an object-centred model of three dimensional shape which is independent of the viewer's observational perspective. This output from the vision-module is then delivered as input to the central cognitive modules which carry out the task of object-categorization.

So, by the account on offer from the functional decomposition camp, the cognitive role of perception is to recover, in the form of internal representations, the properties and relations of an essentially pre-given external environment [31]. Perception is veridical just when the representational description delivered to 'cognition-central' by the perception module accurately captures those properties and relations. On this view, perception and action can be analysed and studied independently of each other, because they are held to be temporally distinct and conceptually separate functions of an intelligent control-system [7]. But however influential such a set of explanatory principles has been, there are good reasons to think that they do not present the most profitable way to proceed. In short, it is a confusion to think that particular perceptual processes can be studied independently of the specific adaptive activities in which those processes are embedded.

3 Active Perception

Let's start by taking a look at visual perception. The view adopted in mainstream A.I. has resulted in most computational vision research being orientated towards static scene analysis, rather than ongoing activity in a world. Even those researchers who have employed more 'dynamic,' ecological concepts, such as optic flow, have done so with the aim of constructing more useful world-models for cognition-central, rather than in the context of specific environmentally embedded activities. But now consider John Haugeland's fridge [15].

[What's] remarkable about our refrigerator aptitudes is not just, or even mainly, that we can visually identify

- the intimate interlinking of perception and action through close sensory-motor couplings between agent and environment,
- the role of the environment in explaining perceptual capacities,
- the use of specialized perception-action modules in robots, and the hypothesized existence of such modules in animals,
- the attainment of overall competence in certain perceptual abilities, even given far-from-optimal performance in relevant sub-competences,
- the part played in perception by sensors with dynamic properties,
- the nature of representation in active perception.

Action and Interaction

In a process of ongoing interaction with an environment, temporal constraints are hardly ever arbitrary. Lurking behind the sense-model-plan-act methodology is a crucial premiss to the effect that, even given accuracy problems resulting from noisy or drifting sensory-motor mechanisms, it is still *possible* to build an adequate, stored world-model, and to manipulate that model in real-time. This is required so that, for the purposes of planning action, operating in an actual world can be ignored in favour of the internal representations. But, as adaptive behaviour researchers (and others) have often observed, in scenarios where an autonomous agent’s domain of activity is a dynamically changing and/or uncertain environment, a commitment to the necessity of maintaining an accurate internal world-model could well be a devastating error. Due to an explosion in the demands placed on representational and computational resources, the problem becomes intractable on the time-scales relevant to the realization of adaptive behaviour (a fact which would signal the untimely end of many a predator-threatened animal).

But are these difficulties products of the way in which perception and action are conceptualized in architectures committed to decomposition by function? Evidence that the answer to this question may well be “yes” can be gleaned from a consideration of animat control systems in which perception and action are intimately intertwined through the realization of tight feedback loops between the agent’s sensory-motor mechanisms and the environment. For example, Franceschini *et al.* [12, 13] describe how real-time visual guidance of a path-finding autonomous mobile robot, demonstrating obstacle avoidance as one self-contained competence, was achieved through the speeds of its drive and steering motors being adjusted via purely local visual feedback loops. This close sensory-motor coupling enables the fully-developed robot to function in certain classes of environments in which it had not been debugged, and to succeed not only at path-finding tasks involving stationary spatially located goals (for which the hardwired control system was specifically designed) but also at similar tasks involving non-stationary goals (for which the hardwired control system was not specifically designed). This ‘natural adaptability’ of the control system rests on the intimate link between perception and action. Indeed it seems that — within certain restrictions on how fast a moving goal is travelling in relation to the robot’s sensory-motor capacities — *any* goal-finding architecture featuring close sensory-motor couplings will have a natural adaptability to transfer from environments featuring a static goal to those featuring dynamic goals. (This adaptability will fail if, in the static-goal case, the sensing mechanism has been picking up not properties of the goal, but properties of

some other fixed object in sensory range — such as a window — which had a fixed relation to the stationary goal).⁴

Franceschini *et al.*'s robot (more on which below) is an example of a *behaviour-based* control architecture, as pioneered by Brooks and his colleagues.⁵ The 'behaviour-based' approach advocates architectures with no central reasoning systems and no manipulable symbolic representations. Given the intuition that a situated agent should operate by continuously referring to its sensors as opposed to some internal representation, the process of attempting to build a centrally stored, objective world model is rejected as constituting a positive hindrance to real-time activity in a messy environment. Individual behaviour-producing systems, called 'layers', are designed to be individually capable of — and to be generally responsible for — connecting the robot's sensing and motor-activity in the context of, and in order to achieve, some ecologically relevant behaviour. So each layer is closely coupled to the robot's environment along what might be called a 'channel of ecological significance.' Starting with layers which achieve simpler behaviours (such as 'avoid objects' and 'explore'), layers are added, one at a time, to a debugged, working robot, so that overall behavioural competence increases incrementally. The layers run in parallel, affecting each other only by means of suppression or inhibition mechanisms.

Any account of perception based on tight sensory-motor couplings effectively makes a prediction that, in order to achieve ongoing perceptually guided activity, autonomous

built in, because it implicitly provides ‘recognition’ of the correct signal through the failure of the system with any other signal.”

Moreover there is no need to hypothesize two separate processes — one to recognize the song, and one to song,

whose eyes, necks and bodies are all movable, saccades (which transfer the image of an

one can, if one wishes, readily treat behaviour as an effect of the milieu. But

that rejection.¹⁰

Before confronting that issue of cospecification head-on, we need to place the theory of affordances in the adaptive behaviour framework assumed throughout this paper. So the intrinsic ‘good’ or ‘ill’ of an affordance would generally need to be ‘cashed out’ in terms of positive or negative survival and reproduction prospects for the agent. Consider an adaptive problem such as catching some prey. What will be detectable (via sensory information) will be environmental objects, events, or situations which Miller and Freyd [25] call ‘fitness affordances.’ Just like their Gibsonian ancestors, Miller and Freyd’s evolutionarily defined fitness affordances ‘point both ways.’ A prey is only a prey with respect to some predator. It may, itself, be a predator to some other creature on which *it* preys. Consequently, fitness affordances are irreducibly cospecificationary, whilst being perceiver-independent in the sense that “fitness effects are imposed by natural selection whether the organism likes it or not; they cannot be eliminated through subjective denial or wishful thinking” (Miller and Freyd, 1993, p.16). This modification, whilst significant, is clearly in the spirit of Gibson’s approach.

But now how can we explicate ‘organism-environment cospecification’? Here we can turn to a different conceptual language — that of dynamical systems theory¹¹ — to provide a way of conceptualizing organism-environment relations in which this somewhat vague concept becomes much more concrete.

A nervous system is a complex dynamical neural network that constitutes the basis of the control system for a situated agent. (In general, the agent’s sensory-motor mechanisms should be thought of as part of its control system.) This control system is embedded in a continually changing physical medium, with which it interacts in such a way that the network’s intrinsic dynamics are regularly perturbed as a result of the agent’s sensory-motor activity. The agent’s control system and the medium in which that control system operates can be conceptualized as two *coupled* dynamical systems. Two theoretically separable dynamical systems are said to be coupled when they are bound together in a mathematically describable way, such that, at any particular moment, the state of either system fixes the dynamics of the other system; that is, each system fixes the principles governing change in the other system. In formal terms, this means that some of the parameters of each system either become, or become functions of, some of the state variables of the other.¹²

The situation in which one system fixes the dynamics of another system through coupling is not to be equated with a relation according to which one system specifies the state of a second. Each system biases the intrinsic possibilities for change already present in the other. The relation is one of influence of dynamics rather than specification of static state. If we begin by thinking of an animal nervous system as a non-coupled dynamical system, then we can conceptualize its intrinsic dynamics as generating a space of possible perturbations which the system can undergo as a result of coupling to a physical medium. Then, through sensory-motor activity, the dynamics of an animal’s nervous system are continually perturbed in accordance with the adaptive couplings ‘discovered’ by evolution.¹³

In effect, we have now explicated the mechanism of attunement. It is the discovery, by Darwinian natural selection, of adaptive sensory-motor couplings. This evolutionary pro-

¹⁰Varela *et al.* make similar criticisms of the Gibsonian framework (Varela *et al.*, 1991, pp.202-5).

¹¹Abraham and Shaw offer a friendly, but thorough, introduction to dynamical systems theory [1].

¹²For

cess operates on the nervous system of the organism. So, given the sense of ‘information’ relevant to the Gibsonian framework, the mechanisms resulting in attunement are not informationally sensitive, and so cannot be tuned to informational properties. This may look as if we are further away from organism-environment cospecification. But notice that our consideration of adaptive attunement has confined itself to a level of explanation characterized by physical processes of stimulus and response, occurring in the networks making up the sensory-motor and nervous systems. That is not the place to locate meaningful ecological ‘objects’ such as fitness affordances. So where should one look?¹⁴

The dynamical systems perspective provides a framework in which the dividing line between the ‘internal’ states of the agent and the ‘external’ states of the environment is something to be imposed in context, relative to the interests of the observer. This is because whilst it is useful, under certain circumstances, to think of *agent and environment*, (and not merely sensory-motor control system and physical medium) as separate, but coupled, dynamical systems, it is equally valid to redescribe the coupled agent-environment system as one larger dynamical system, in which the observed patterns of interaction between the agent-system and the environment-system are properties of that larger system [4]. In fact, it is perfectly legitimate to think of the one agent-environment system as *primary*, and to consider any analysis in terms of coupled dynamical systems as an *abstraction*

not, in themselves, meaningful properties of the agent's world. The way in which the complete agent, embedded in its environment, proceeds actively to exploit those stimuli brings forth meaning. It is this enactive process that

of sensory-motor activity. Active perception takes place in meaningful worlds.

Acknowledgments

This work was supported by British Academy major award no. 92/1693. Many thanks to Maggie Boden, Matthew Elton and Inman Harvey for their valuable comments on an earlier version of this paper.

References

- [1] R. H. Abraham and C. D. Shaw. *Dynamics - The Geometry of Behavior 2nd edition*. Addison-Wesley, Redwood City, California, 1992.
- [2] M. A. Arbib. *The Metaphorical Brain 2: Neural Networks and Beyond*. John Wiley and Sons, New York, 1989.
- [3] D. H. Ballard. Animate vision. *Artificial Intelligence*, 48:57–86, 1991.
- [4] R. Beer. A dynamical systems perspective on autonomous agents. Technical Report 92-11, Case Western Reserve University, Cleveland, Ohio, 1992.
- [5] M. A. Boden. Autonomy and artificiality. To be published in *A.I.S.B. Quarterly* and D. Cliff (ed.) *Evolutionary Robotics and Artificial Life* (provisional title), forthcoming.
- [6] R. Brooks. A robust layered control system for a mobile robot. A.i. memo 864, M.I.T. A.I. Lab, 1985.
- [7] R. Brooks. Intelligence without reason. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, pages 569–95, San Mateo, California, 1991. Morgan Kaufman.
- [8] R. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–59, 1991.
- [9] R. A. Brooks and L. A. Stein. Building brains for bodies. Artificial Intelligence Lab Report 1439, M.I.T., 1993.
- [10] D. Cliff, I. Harvey, and P. Husbands. Explorations in evolutionary robotics. *Adaptive Behavior*, 2:73–110, 1993.
- [11] T. S. Collett and M. F. Land. Visual control of flight behaviour in the hoverfly, *syritta pipiens l.*

- [15] J. Haugeland. Mind embodied and embedded. Draft of paper to appear in the proceedings of 'Mind and Cognition: an International Symposium'. Institute of European and American Studies, Academia Sinica, Taipei, Taiwan., 1993.
- [16] I. D. Horswill and R. A. Brooks. Situated vision in a dynamic world: Chasing objects. In *Proceedings of the Seventh Annual Meeting of the American Association for Artificial Intelligence*, pages 796–800, St. Paul Minnesota, 1988.
- [17] P. Husbands, I. Harvey, and D. Cliff. Circle in the round: State space attractors for evolved sighted robots. *Proceedings of the N.A.T.O. Advanced Study Institute Workshop on the Biology and Technology of Intelligent Autonomous Agents*, forthcoming.
- [18] H. Komatsu and R. H. Wurtz. Relation of cortical areas mt and mst to pursuit eye movements iii: interaction with full-field stimulation. *Journal of Neurophysiology*, 60:621–44, 1988.
- [19] M. F. Land. Similarities in the visual behavior of arthropods and men.

- [29] T. Smithers. On behaviour as dissipative structures in agent-environment system interaction spaces. Faculty of Information, University of San Sebastian, 1993.
- [30] T. J. van Gelder. What might cognition be if not computation? Technical Report 75, Indiana University, Cognitive Sciences, 1992.
- [31] F. J. Varela, E. Thompson, and E. Rosch. *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press, Cambridge, Massachusetts, and London, England, 1991.
- [32] B. Webb. Modeling biological behaviour or ‘dumb animals and stupid robots’. In *Proceedings of the Second European Conference on Artificial Life*, pages 1090–103, 1993.
- [33] M. Wheeler. For whom the bell tolls: The roles of representation and computation in the study of situated agents. Cognitive Science Research Paper 320, University of Sussex, 1994.
- [34] M. Wheeler. From activation to activity: Representation, computation, and the dynamics of neural network control systems. *A.I.S.B. Quarterly*, 87:36–42, 1994.
- [35] S. W. Wilson. Knowledge growth in an artificial animal. In J. J. Grefenstette, editor, *Proceedings of an International Conference on Genetic Algorithms and their Applications*, pages 16–23, Pittsburg, PA and Hillsdale, New Jersey, 1985. Lawrence Erlbaum Associates.
- [36] B. M. Yamauchi and R. D. Beer. Sequential behavior and learning in evolved dynamical neural networks. Available as Case Western Reserve University Tech. Report no. CES-93-25. Cleveland, Ohio, Submitted to *Adaptive Behavior*.