

In Defence of Functional Analysis

Joe Faith

School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK
josephf@cogs.susx.ac.uk

Abstract

Computationalism presupposed a modular-functional analysis of cognitive behaviour, and its failure has encouraged the search for alternative analytical techniques, such as behavioural decomposition and dynamical systems theory. This paper argues that these alternatives can be no more than useful heuristics, and that a functional analysis is necessary for an understanding of intentional adaptive behaviour, but that this need not imply the existence of cognitive modules.

1 Mechanism, Explanation and Elimination

The function of most of the interactions between a living thing and its environment are directly metabolic — digestion, respiration, photosynthesis, etc. In the case of the simplest of organisms these are the *only* interactions engaged in. However, more complicated organisms also engage in interactions in which the function is not directly metabolic, but depend upon an objective relationship between that bit of the environment with which it is engaged, and some other bit that may hold some adaptive benefit — what Gibson described as an *affordance* [27]. For example, when an animal reacts to light hitting its retina it does so, not because this reaction in itself helps the animals survival, but because the pattern of light on the retina stands in some objective relationship to some feature of the environment that may be. (In some cases — such as plant phototaxis — the two functions are combined in one behaviour). In other words the function of these interactions is *about* something else in the environment. These interactions are thus necessarily *intentional*: they can only be understood by reference to the distal object. If the organism is to be consistently adaptive then the objective relationship between it and that affordance must be reflected¹ in some way in the organism’s nervous system (if not in its psychological experience).

If a behaviour is adaptive in this way then it must display some form of regularity: if there is food available, then the organism must reliably find it; if there is

a predator, then it must avoid it; and so on. In general, adaptive behaviour presupposes behaviour that is regular, reliable or robust with respect to some *affordance* offered by the environment [27]. Obviously not all such regular behaviours will be adaptive — consider an organism that reliably runs *towards* a predator — nor does the regularity have to be perfect. Nonetheless, behavioural regularity is a pre-requisite of behavioural adaptivity.

How can a mechanism produce such behavioural regularities — what is the nature of this “reflection”? One way is for them to be codified in a set of rules that use representations of the state of the environment (such as “if there is some food, then move towards it”), and then to build a system that follows these rules. This is the intuition underlying computationalism, with all its attendant problems [10].

An alternative is described by Daniel Dennett:

But how *could* the order be there, so visible amidst the noise, if it were not the direct outline of a concrete orderly process in the background? Well, it *could* be there thanks to the statistical effect of very many concrete minutiae producing, as if by a hidden hand, an approximation of the “ideal” order. Philosophers have tended to ignore a variety of regularity intermediate between the regularities of planets and other objects “obeying” the laws of physics and the regularities of rule-following (that is, rule-*consulting*) systems. These intermediate regularities are those which are preserved under selection pressure. [18, p43]

Thus behavioural regularities need not be produced by a corresponding functional entity in the underlying mechanism — a “controlled variable” [48] — rather they can be a collective by-product of many essentially independent processes. Douglas Hofstadter [31, p642] gives the example of a computer system whose performance degrades disastrously when there are more than, say, 35 users on the system. A naive observer may suspect that there is a mechanism that detects the number of users and controls performance accordingly. In fact the drop in performance is due to the system “thrashing” — spending so much time swapping between users, that there is no time left for useful processing. This regular, environmentally-contingent, behaviour is a result

¹ “Reflected” is used, not in the mirror-like sense, but in the sense of “the change in the weather was reflected in people’s dress”.

of what Hendriks-Jansen calls “interactive emergence” [30] between the mechanism and its environment, and is not caused by any discrete functional entity in the former. These two ways of producing behavioural regularities correspond to Marr’s distinction between Type I and Type II mechanisms [37].

A rule-following, Type I, system has at least two levels of organisation: the lower level of the underlying mechanism, and the higher level of the rules that it follows. For example, a digital computer is realised in an electro-dynamical system that displays higher-level organisation in the form of the virtual machinery of its software. This higher level of organisation is as real as the lower level, and cannot be reduced to it: the higher level can be *explained* in terms of the lower, but cannot be *explained away*. A full explanation of the system must describe both levels. A Type II system, on the other hand, has no such higher level of organisation. Its behaviour when interacting with an environment may possibly be describable in terms of strict rules, but these rules are purely a property of its behaviour and do not correspond to any entities in (or level of organisation of) the underlying mechanism.

The regularities in adaptive, behaviour are usually described using *intentional* terminology. For example: “the agent turned left because it thought that there was some food over there and it was hungry”. If the behaviour were produced by a Type I mechanism that instantiates these rules as a level of organisation, then these descriptions would be true explanations — they would not just describe the behaviour, but also correctly identify features of the mechanism that *produced* that behaviour. In our example, there would be some functional entity in the organism’s mechanism that corresponds to the sensing of the position of the food on the various different occasions when the intentional description could be applied.

However, there need be no such correspondence. If the mechanism were of Type II then the intentional interpretation may accurately describe the *behaviour* of the agent even though it bears no relationship to the way in which the underlying mechanism *produces* that behaviour. If this is the case then we would be justified in eliminating the intentional explanation in favour of an analysis purely in terms of the underlying mechanism, understood as a dynamical system [45][12]. The intentional interpretation would be descriptive, but not explanatory.

Consider this analogy. Suppose that a ball were affected by two forces: one pushing north, and one west. We can describe the behaviour of the ball as being due to a single force acting north-west, and this is a perfectly accurate, and predictive, description. However the entities postulated by this description bear no relation to the mechanism that produced the behaviour. It is thus equivalent to an intentional interpretation of a Type II system. A true *explanation* of the behaviour of

the ball would eliminate the resultant force in favour of the underlying components, just as the intentional interpretation should be eliminated in favour of a dynamical systems explanation. Note that we cannot distinguish between the descriptions of the behaviour on a purely empirical basis, rather we have to investigate the mechanism that produced the behaviour.

The central problem of naturalising intentionality is in what sense is it true that an agent interacts with a distal object, rather than with proximal stimuli. This paper avoids this problem, and instead asks in what sense is it *useful* to make such claims. The conclusion is that if an ascription of intentionality to an agent is to have explanatory, rather than just descriptive, virtue, then it must be grounded in the functional organisation of the mechanism underlying its behaviour.

2 The Evolution of Adaptive Behaviour

There is no *a priori* reason why natural or artificial adaptive systems should be of either type: both are capable of displaying exactly the same behaviours. However there are practical reasons why Type I mechanisms will be favoured when behaviours emerge through a process of learning, development, or evolution — unlike Hofstadter’s example, which was an accidental by-product of a design produced for another purpose.

Recall that adaptive behaviour demands that an agent coordinate its activity with respect to an affordance that the environment offers. To do this the agent must detect it — where “detect” is used in the widest sense of “being influenced by its presence”. Given some affordances, and some sensory mechanisms, this can be done very simply. Consider one of the vehicles discussed by Braitenberg [6]: a mobile robot with two driven wheels, a castor to prevent toppling, and two light sensors, both pointing forwards with one to the left and one to the right. If the output from the left sensor is connected to the right motor, and *vice versa*, then the robot will continuously turn and move in the direction of the strongest light source. This reliable behaviour is not due to any internal representations or a rule-following mechanism. The presence of an affordance can also be detected in much more complex sensory input, given a suitable sensory mechanism — such as the tracheal tubes of the female cricket, which are precisely tuned to detect the male’s chirp [52][32].

In these cases there is a very simple relationship between the presence of an affordance and a particular sensory input, and a Type II mechanism is capable of using this relationship to produce adaptive behaviour. However organisms often have to use general-purpose mechanisms (such as retinal arrays or olfactory bulbs) rather than specialised detectors (such as tracheal tubes, or pheromone detectors). As the sensory mechanisms become more general, so the presence of particular affordances become less explicit. For example, in olfactory

systems that detect pheromones, the presence of the affordance that the pheromone signals is proportional to the activity of a local set of receptor cells. In contrast, when an odour is detected by an olfactory bulb it produces an oscillatory pattern that is distributed across a very wide area [44]. How can neural mechanisms evolve, or learn, to use such general sensory systems to reliably detect the presence of an affordance?

In such cases, the more successful mechanisms will be those that are better able to generalise from one case of a behavioural regularity to another. Consider shift invariance. In this case the regularity required is a similar response to an image that may be presented at different positions on a retina. It would be possible to teach, or evolve, a Type II mechanism to respond suitably to a number of particular image presentations. However, since it is a Type II mechanism then the way in which it achieves these responses may have nothing in common with each other: the behavioural regularity may not be due to a regularity in the underlying mechanism.

This mechanism is likely to be robust against certain changes in input: the addition of noise to the input would only change each input channel by a small amount, and thus we may expect the response of the mechanism to be similar. However other changes, such as moving the image on the retina, result in each input channel changing by a large amount — spots that were dark will now be light, and *vice versa*. There is no reason why a previously successful mechanism would produce similar behaviour given such different novel inputs. By contrast, a Type I solution would, by definition, involve some functional element that responds similarly to the same pattern regardless of its particular position. Such a solution, if it can be found, will generalise more robustly and will thus be favoured by evolution or a learning regime.

For example, one possible (though completely artificial) Type II solution to learning shift invariance would be a look-up table that lists a number of different retinal inputs and their required responses. This solution is completely incapable of generalising beyond the listed inputs, unlike a Type I solution that depended on noticing what the positive cases had in common.

Type I organisation need not be a property of the internal mechanism; it may also be achieved through active perception. For example *Drosophila* solves shift invariance, not through a detector that can recognise a pattern anywhere on the retina, but by moving its eye until a template is matched [19]. A similar strategy is used in [43] to solve the problem of distinguishing small objects from large ones using only a one dimensional array of proximity sensors. If the vehicle “wall-hugs” any object that it comes across, then the size of the object being hugged will be inversely proportional to the rate of turning, which can be detected as a difference in left and right wheel speeds. Both methods achieve affordance-

behaviour regularities by unifying affordance-detection through a single functional unit — whether this unification is done internally or through the agents own activity.

Thus we would expect Type I affordance-detection to be an almost inevitable product of even the most unbiased learning or evolution regimes [49]; and, indeed, this is the case. In [22] Floreano and Mondada describe the artificial evolution of a neural network controller for a Khepera-style robot. Its task is to explore a simple arena, returning to a recharging “base” that is demarcated by a black floor patch, directed by a bright light. The fitness of an individual rises with the total amount of movement, along with avoiding obstacles. Although the need to recharge is not an explicit element in the fitness function, it is a behaviour that obviously must be evolved in order to maximise fitness. Thus the location of the base is a feature of the environment that it would be very useful for the robot to detect (though it could be found by a random search, albeit less efficiently). The most fit individual was found to be using one hidden node of the network in order to do this — its activation corresponded to the distance from the base, reaching a maximum when it was “home”. As the authors note:

In this experience the robot autonomously evolved the ability to use the raw sensor data and built an internal representation of the world in order to find the recharging area and return to this place at a given time. This behaviour is based on an accurate evaluation of the battery residual time and on an internal representation of the environment. In fact some of the hidden nodes displayed activation levels that clearly mapped the environment geometry. [41]

The second example is [29], in which a team from the University of Sussex evolved a neural network to control the movement of a camera-head mounted on a gantry, whose motion is designed to mimic that of a wheeled robot. In this case both the internal network, and the morphology of the visual sensors, was available for selection. The task presented was to approach a white triangular target, whilst avoiding a rectangular one. The successful robot used two sensors, one with a visual field above the other. It locates the triangle by rotating on the spot until just the lower sensor sees white, when it moves straight ahead. This has the effect of fixating the robot on the oblique edge of the triangle. As the triangle looms up such that both sensors go high, or if the motion causes the edge to be lost, then the robot will start to rotate until the edge can be fixated again. The rotate/move-straight distinction is effected by a single unit that takes an inhibitory connection from the upper sensor and an excitatory link from the lower, and is thus only fully activated when the robot is facing towards the triangle’s edge. Therefore this robot uses a mixture of

active perception, a sensor morphology closely tied to the structure of the environment, and a representational architecture in order to produce adaptive behaviour.

The evolutionary pressure to generalise behaviour thus produces a pressure to localise function. Thus we find that the most robust, stereotyped, behaviours are produced by the most functionally specialised mechanisms. For example, sensory-motor behaviours, such as saccades or fixation movements, that have to be very fast and reliable tend to be produced via very clearly defined topographic cortical maps [39].

It is also worth remembering that not all components are localised, structurally individuated entities like hearts and lungs: in small animals the functions of respiration and circulation can be achieved by diffusion processes. This does not imply that there are no entities — stomata etc — that carry this function, but that they form a distributed, functionally individuated, “component” or subsystem; rather than a localised, structurally individuated, one. Similarly, componential-functional decomposition in neural networks need not imply the existence of “grandmother” cells, or even of clearly delimited modules. Something can play a well-defined functional role with respect to other components even if it is not topographically localised. For an intuition pump, think of the geographically diffuse functional components of human societies, such as political organisations, classes, companies etc. As was shown above, it is usually only the most stereotyped behaviours that result in localised functional modularisation; and even these do not work in isolation. For example, although only a small number of neurons are directly involved in the gill-withdrawal reflex of *Aplysia*, up to 300 others are simultaneously activated, since siphon stimulation *also* causes many other behaviours: mantle contraction, inking, mucus release, postural changes, respiratory pumping etc. [2]

Functions Without Computations

It must be emphasised that a functional decomposition is relative to a behaviour of the overall system [14]. This is the crucial difference between this approach to functional analysis, and the *modular*-functional analysis of classical computationalism. Fodor [23] argues for the existence of modules that are the prior explanatory atoms of all cognitive behaviour. Modules are general purpose, they play the same role in all behaviour, and so have a fixed function. However, the functional analysis given above *starts* from a particular behaviour, and then asks how it is achieved. Analysing different behaviours may reveal a different functional decomposition with no component playing the same role in each case. It is modularity, rather than functional analysis *per se* that defines a computationalist perspective.

When a single unit is described as representing the presence of a target triangle, this is a description of how

a mechanism achieves a behaviour. Representation is not what an entity *is*, but rather what it *does* in a behavioural context. Representation is not merely a correlation between internal and external state that only exists for an external observer, but a relational, functional property between a mechanism and a particular behaviour that it displays.

3 Functional Analysis and Dynamical Systems Theory

What do we gain by a functional analysis of such simple systems? After all, if a functional analysis is needed in a simple case like the triangle-seeking robot, then there seems no principled reason why it could not be applied in even simpler cases. For example, the neural mechanism of the Braitenberg light-seeking vehicle is actually no more than a pair of crossed wires. However we *could* describe it in functional terms as an input module (comprising the two sensors) that passes a representation of the world (the state of the two wires) to the output module (the motors). Although no principled “bottom line” for functional explanation can be given, this need not invalidate its use. After all, if we consider a vacuum that contains just a few molecules of a gas, then concepts such as “temperature” and “pressure” would seem superfluous compared to a far more precise kinematic description of the molecular motion. However this does not mean that the bulk gas properties are not well-defined; it is only when the number of molecules increases that the explanatory power of these properties becomes essential.

A more principled objection is that functional-intentional explanations should be eliminated in favour of a description of the agent in its environment as a coupled dynamical system (DS) [28][50]. At this point two versions of the DS Theory should be distinguished, depending on what are taken as the state variables of the system. The state variables are the explanatory “atoms” of DST. Once a set of state variables has been identified, then their relationships and effects on each other are described with a set of evolution equations. Factors that are external to, but impact upon, the system are included as parameters to those equations. Two dynamical systems in which the state variables of one system act as the parameters for the other (and *vice versa*) are described as being *coupled*, and together form a new super-system which may, in turn have its own environment and parameters. A dynamical system is analysed by examining the structure and topology of the phase space of the system, and any basins of attraction that underpin any identifiable behavioural modes — what Agre and Chapman call “routines” [1].

The first, weaker, version of DST allows mental and intentional attributes to be used as state variables — such as the motivations, beliefs and decisions modelled in [11]. This approach thus shares the fundamental Cartesian

assumption of classical computational psychology, that cognition takes place in a mental “space” built from representations of the world. The only difference is that the rules for the manipulation of these representations are essentially temporal. The weak form of DST therefore stands in the same relation to computationalism as classical connectionism [16]: it shares an explanatory framework, but uses more complicated rules.

The stronger form of DST therefore restricts state variables to be non-mental, i.e. *physical*, properties of the system and its environment [4] [46]. This approach was first used 50 years earlier by Ross Ashby [3]. In the case of agents built from a neural network, the obvious choice of state variables are the activation of single neurons or neural masses [33][24]. However, this restriction has consequences for how DST can be used.

DST is supposed to provide an explanation of how cognitive behaviour is produced, rather than simply describing the dynamics of the internal mechanism. In order to do this a DST model must therefore include not just a whole agent but also its *environment* (though see [34]). Let us assume for the moment that determining the evolution equations for state variables that are internal to the agent is unproblematic. The environment can then be handled in one of three ways. The first is to ignore it by leaving any environmental impact on the system as undetermined parameters: this is obviously no answer for analysing the behaviour of whole agents. The second, as advocated in [46] and [4], is to treat the environment as a dynamical system in its own right, tightly coupled to that of the agent. This, however, is a form of Laplacean reductionism. Consider trying to produce a DST model of a bird trying to land in a tree. The bird must coordinate its body with a swaying branch using its eyes, brain, and muscles. The obvious state variables for the bird will include retinal cell inputs, the activations of the neurons in the visuo-motor system, and muscle nerve outputs. However the retinal inputs to the system will be affected by the most trivial changes in the environment: the wind catching a leaf and causing a shadow to move, for instance. This perturbation of input will, in turn, alter the trajectory in the phase space of the system. If we want to know how the bird lands reliably despite the incidental movement of so many factors in the environment, then it seems as though we will have to model the tree in as much detail as the nervous system of the bird. This is impossible in practise, even if we agree with Laplace that it may be possible in principle. The only case of a full DST analysis of a whole agent-environment system that I am aware of is [33], in which the environment is completely static. In [4] the neural network controller for a hexapod robot is modelled as a dynamical system, but its environment is treated as the body that it controls (the environment external to the body is again assumed to be static), and only a localised 5-neuron subsystem

is *analysed* as such. (This is not meant to detract from the great subtlety of the evolved design. The point is that it is not possible to fully appreciate it from a purely dynamical systems perspective.)

The only alternative for DST is to postulate features of the phase space of the agent-environment system that are immune to the incidental, un-modellable, changes to state variables. This is the way in which Walter Freeman uses DST to model oscillations in the olfactory bulb [25]. Two points should be noted about this model. The first is that it is *not* an attempt to model a whole agent-environment system, rather it is a model of an isolated functional module with a well-defined input and output. Second, the whole point of this model and the experimental *in vivo* work on which it is based [51], is to show how certain characteristics of the phase space of the olfactory bulb are immune to certain changes in input: in particular that the spatial amplitude pattern of the dominant oscillations carries odorant information [26], despite the stochasticity of epithelial receptor activity [35]. In other words, DST is being used to examine how the olfactory bulb fulfills the function of the classification of odorant information; in this case it underpins, rather than undermines, a functional analysis [21].

In general, DST tries to relate particular behaviours to topological features in the phase space of the system. This requires two things. The first is that certain state variables of the agent and environment have to be ignored at certain times: for example if we want to know how an agent is physically negotiating an obstacle, then the olfactory input will not be relevant. This corresponds to taking projections of the phase space in order to reveal topographic regularities. A system could well be in a well-defined limit cycle with respect to 2 state variables, whilst another varies seemingly randomly. The underlying order is only exposed if we consider a 2-variable sensory-motor subsystem, for which the state variables of the larger system act as parameters. Changes to these parameters may well drastically alter the topology of the subspace causing a change in the behavioural mode, and hence they cannot be ignored but, until they do change, the exhibited behaviour must be analysed with respect to a subset of state variables. If we recall that state variables measure physical properties of parts of the system, then dynamical systems analysis requires that a particular behaviour is related to properties of a physical subsystem. Second, the topographic features to which behavioural modes are related are identified with respect to *values* of state variables — i.e., with respect to the internal state of a subsystem. The internal state may not be a *static value*, but rather — as in the case of the olfactory bulb — must be a *stable mode* of activity.

Therefore, contrary to the many claims made for it, if DST is to be used to analyse the adaptive behaviour of whole agents in dynamic environments it will require

that the presence of affordances in the environment be related to the functional role of the state of subcomponents of the system. In other words, DST in practise produces a functional, componential, and hence *intentional*, analysis. Adaptive systems, like everything else, may be *modelled* as a dynamical system. The point is that they cannot be *understood* as such.

4 Functional Analysis and Behavioural Decomposition

Another proposed alternative to the componential-functional analysis of adaptive behaviour is behavioural decomposition. Instead of decomposition into simpler *functions*, the overall behaviour should be decomposed into simpler *behavioural capacities*. As Cummins puts it:

A cook's capacity to bake a cake analyzes into other capacities of the "whole cook". . . . My capacity to multiply 27 time 32 analyzes into the capacity to multiply 2 times 7, to add 5 and 1, etc. These capacities are not (so far as is known) capacities of my components: indeed, this analysis seems to put no constraints at all on my componential analysis. [15, p29]

Brooks [7][8] has demonstrated how each of these whole-agent capacities may be achieved by a "layer" in a control architecture, connected to the sensors and motors of the system and working semi-independently of each other. (Also see [36] [30] [47].) Contrast this to the normal modular decomposition [23] in which only the sensory and motor modules are connected to the outside world, with all other modules communicating between themselves using representations.

However, given a behavioural decomposition and a corresponding layered control architecture, the problem still remains of *how* each of these layers achieve their, admittedly simpler, intentional behaviours. Brooks *et al* argue that this can be achieved "by using the world as its own best model" — in other words without any complex transformations of sensory input. However, unless "magic sensors" are available which detect the presence of an affordance directly (and pheromone detectors are a common natural example), then this will not be possible. Even the most cited implementation of Brooks' architecture, Mataric's *Toto* [38], uses clearly defined, localised components to indicate states of the agent-environment interaction. As Brooks states:

My earlier paper [10] is often criticised for advocating absolutely no representation of the world within a behaviour-based robot. This criticism is invalid. I make it clear in the paper that I reject traditional Artificial Intelligence representation schemes. I also made it clear that I reject

explicit representations of goals within the machine.

There can, however, be representations which are partial models of the world — in fact I mentioned that "individual layers extract only those *aspects* of the world which they find relevant — projections of a representations into a simple subspace". The form these representations take, within the context of the computational model we are using, will depend on the particular task those representations are to be used for. [9, p19]

Behavioural decomposition also often implicitly assumes that adaptive agents have a discrete set of distinct units of behaviour [20] — one for each layer. For example, Hendriks-Jansen describes species-specific fixed action patterns as "natural kinds" — atoms out of which the overall behavioural repertoire is built [30]. The same intuition underlies the use of extrinsic fitness functions to evolve neural network controllers for adaptive agents, scoring each member of the population on their performance on a precisely defined behaviour. This is closely related to Dawkins' genetic determinism, in which the "extended" phenotype of the organism — which includes behavioural properties — is divided into discrete traits, each of which is encoded by a single gene [17].

Beer and Gallagher [5] have tried to avoid this behavioural atomism by advocating the use of *intrinsic* fitness functions, in which the selection is more "natural". Instead of reproductive success being determined by success in a behavioural trial, the members of the population have autonomous metabolic and reproductive cycles. This means that, as in natural selection, the evolved agents are judged on their overall "way of life" rather than performance over a set of atomistic behaviours. (See [42] for a simple example.)

5 Conclusion

If we want to understand how evolved mechanisms produce adaptive behaviour, then there is no philosophical or practical alternative to functional analysis. A dynamical systems analysis, though a useful counterweight to computationalist assumptions, is based on a Laplacean philosophical error and reduces to functional analysis in practise. Behavioural decomposition is a useful heuristic for synthesis, but assumes the existence of behavioural atoms, and begs the question of how those atoms are achieved.

A functional analysis of the adaptive behaviour of an agent must relate how it is coordinated with respect to affordances offered by the environment, to internal functional entities. Therefore the analysis will be necessarily *intentional*. Instead of discarding functional, intentional analyses, we should investigate how they can be used in non-modular, non-computational ways.

Acknowledgements

Thanks to Maggie Boden, Ron Chrisley and Jason Noble for useful comments. A version of this paper has been accepted to the fourth European Conference on Artificial Life.

References

- [1] P.E. Agre and D. Chapman. *Pengi: An implementation of a theory of activity*, pages 196–201. 1987.
- [2] J.S. Altman and J. Kien. Highlighting *aplysia*'s networks. *Trends in Neuroscience*, 13(3), 1990.
- [3] W.R. Ashby. *Design for a brain: The origin of adaptive behaviour*. Chapman Hall, 1952.
- [4] R.D. Beer. A dynamical systems perspective on autonomous agents. Technical Report 92-11, Case Western Reserve University, Cleveland, Ohio, 1992.
- [5] R.D. Beer and J.C. Gallagher. Evolving dynamical neural networks for adaptive behaviour. *Adaptive Behaviour*, 1(1), 1992.
- [6] V. Braitenberg. *Vehicles : experiments in synthetic psychology*. MIT Press, 1984.
- [7] R. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, (2):14–23, April 1986.
- [8] R. Brooks. Challenges for complete creature architectures. In Meyer and Wilson [40], pages 434–443.
- [9] R. Brooks. Intelligence without reason. Technical Report 1293, MIT Artificial Intelligence Laboratory, April 1991.
- [10] R. Brooks. Intelligence without representation. *Artificial Intelligence*, (47):139–159, 1991.
- [11] J.R. Busemeyer and J.T. Townsend. Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, (100):432–459, 1993.
- [12] P.M. Churchland. Eliminative materialism and the propositional attitudes. *Journal of Philosophy*, (78), 1981.
- [13] D. Cliff, P. Husbands, J.A. Meyer, and S.W. Wilson, editors. *From animals to animats 3: Proceedings of the third international conference on the simulation of adaptive behaviour*. MIT Press, 1994.
- [14] R. Cummins. Functional analysis. *Journal of Philosophy*, 72:741–765, 1975.
- [15] R. Cummins. *The nature of psychological explanation*. MIT Press, 1983.
- [16] R. Cummins and G. Schwarz. Radical connectionism. *The Southern Journal of Philosophy*, XXVI - supplement:43–72, 1987.
- [17] R. Dawkins. *The Extended Phenotype*. W.H. Freeman, 1982.
- [18] D.C. Dennett. Real patterns. *The Journal of Philosophy*, 88(3):27–51, 1991.
- [19] M. Dill, R. Wolf, and M. Heisenberg. Visual-pattern recognition in drosophila involves retinotopic matching. *Nature*, 365(6448):751–753, 1993.
- [20] B. Enc. Units of behaviour. *Philosophy of Science*, (62):523–542, 1995.
- [21] J.E. Faith. The role of oscillations in the olfactory bulb. Master's thesis, School of Cognitive and Computing Sciences, University of Sussex, Brighton, 1995.
- [22] D. Floreano and F. Mondada. Evolution of homing navigation in a real mobile robot. *IEEE transactions on systems, man and cybernetics*, 26(3), 1996.
- [23] J.A. Fodor. *Modularity of mind: an essay on faculty psychology*. MIT Press, 1983.
- [24] W.J. Freeman. *Mass Action in the Nervous System*. Academic Press, New York, 1975.
- [25] W.J. Freeman. Simulation of chaotic EEG patterns with a dynamic model of the olfactory system. *Biological Cybernetics*, 56:139–150, 1987.
- [26] W.J. Freeman and C.A. Skarda. Spatial EEG patterns, non-linear dynamics and perception: the neo-Sherringtonian view. *Brain Research Reviews*, 10:147–175, 1985.
- [27] J.J. Gibson. *The Ecological Approach To Visual Perception*. Houghton Mifflin, Boston, Ma, 1979.
- [28] M. Giunti. Dynamical models of cognition. In T. Van Gelder and R. Port, editors, *Mind as motion*. MIT Press, 1995.
- [29] I. Harvey, P. Husbands, and D. Cliff. Seeing the light: artificial evolution, real vision. In Cliff et al. [13], pages 392–401.
- [30] H. Hendriks-Jansen. *Catching ourselves in the act: situated activity, interactive emergence, evolution and human thought*. MIT Press, 1996.
- [31] D.R. Hofstadter. Waking up from the Boolean dream, or subcognition as computation. In *Metamagical Themas*. Penguin, Harmondsworth, Middlesex, 1985.

- [32] F. Huber and J. Thorson. Cricket auditory communication. *Scientific American*, 253(6):46–54, 1985.
- [33] P. Husbands, I. Harvey, and D. Cliff. Circle in the round: State space attractors for evolved sighted robots. *Robotics and Autonomous Systems*, (15), 1995.
- [34] N. Jakobi. Half-baked, *ad hoc* and noisy: Minimal simulations for evolutionary robotics. This volume.
- [35] J.S. Kauer. Contributions of topography and parallel processing to odor coding in the vertebrate olfactory pathway. *Trends in Neuroscience*, 14(2):79–85, 1991.
- [36] P. Maes, editor. *Designing Autonomous Agents: Theory and Practise from Biology to Engineering and Back*. Special Issues of *Robotics and Autonomous Systems*. MIT, 1991.
- [37] D. Marr. Artificial intelligence: A personal view. In M. Boden, editor, *The Philosophy of Artificial Intelligence*. Oxford University Press, 1977.
- [38] M.J. Mataric. Navigating with a rat brain: A neurologically-inspired model for robot spatial representation. In Meyer and Wilson [40], pages 169–175.
- [39] M. Merzenich and J. Kaas. Principles of organization of sensory-perceptual systems in mammals. *Progress in psychobiology and physiological psychology*, (9):1–42, 1980.
- [40] J.A. Meyer and S.W. Wilson, editors. *From animals to animats: Proceedings of the first international conference on the simulation of adaptive behaviour (SAB90)*. MIT Press, 1991.
- [41] F. Mondada and D. Floreano. Evolution of neural control structures: some experiments on mobile robots. *Robotics and Autonomous Systems*, 16:183–195, 1996.
- [42] N.H. Packard. Intrinsic adaption in a simple model for evolution. In C. Langton, editor, *Artificial Life: Proceedings of the workshop on artificial life*, Santa Fe Institute studies in the sciences of complexity. Addison-Wesley, 1988.
- [43] C. Scheier and R. Pfeifer. Classification as sensory-motor coordination: A case study on autonomous agents. In F. Moran, A. Moreno, J.J. Merelo, and P. Chacon, editors, *Advances in artificial life: proceedings of the third European Conference on Artificial Life*, number 929 in Lecture notes in artificial intelligence. Springer, 1995.
- [44] C.A. Skarda and W.J. Freeman. How brains make chaos in order to make sense of the world. *The Behavioral and Brain Sciences*, 10:161–195, 1987.
- [45] T. Smithers. Taking eliminative materialism seriously: A methodology for autonomous systems research. In F.J. Varela and P. Bourgine, editors, *Towards a practise of autonomous systems: proceedings of the first European Conference on Artificial Life*, 1992.
- [46] T. Smithers. What the dynamics of adaptive behaviour and cognition might look like in agent-environment interaction systems. In T. Smithers and A. Moreno, editors, *3rd International Workshop on Artificial Life and Artificial Intelligence, The Role of Dynamics and Representation in Adaptive Behaviour and Cognition*, San Sebastian, Spain, 1994.
- [47] L. Steels. Towards a theory of emergent functionality. In Meyer and Wilson [40].
- [48] L. Steels. The artificial life roots of artificial intelligence. *Artificial Life Journal*, 1, 1994.
- [49] C. Thornton. Brave mobots use representation. Cognitive Science Research Paper 401, School of Cognitive and Computing Sciences, University of Sussex, Brighton, 1995.
- [50] T. Van Gelder. The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, submitted.
- [51] G. Viana Di Prisco and W.J. Freeman. Odour-related bulbar EEG spatial pattern analysis during appetitive conditioning in rabbits. *Behavioural Neuroscience*, 98(5):964–978, 1985.
- [52] B. Webb. Robotic experiments in cricket phonotaxis. In Cliff et al. [13].