

Explaining the Behaviour of Springs, Pendulums, and Cognizers

Michael Wheeler

CSRP 284, June 1993

Cognitive Science Research Paper

Serial No. CSRP 284

The University of Sussex
School of Cognitive and Computing Sciences
Falmer
BRIGHTON BN1 9QH
England, U.K.

this challenge fails. As van Gelder points out, the proposed criterion of causal correlation is insufficient as a characterization of the interactive subtleties of the relationship at issue: the arm angle controls the speed of the engine via the throttle valve; so the arm angle is simultaneously a cause and effect *in the very same interaction*. This prevents it from qualifying as a representation of the engine speed by *the criterion of one-way causal correlation*. In addition, van Gelder stresses that any such correlation between arm angle and engine speed exists only whilst the entire system is at equilibrium. When fluctuations in load or pressure occur, the neat correlation between the two variables no longer exists, yet the governor continues to function admirably; and, moreover, these periods of non-equilibrium are the most important. Therefore, the temporary correlation between arm angle and engine speed cannot play any *key* explanatory role.

If the complex interaction under investigation is beyond the explanatory reach of the representationalist, how should it be described? Van Gelder's answer comes from the language of dynamics. The relationship between the centrifugal governor and the engine (and, thus, between arm angle and engine speed) is one of dynamical coupling. Such a relation obtains when two separable dynamical systems 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. Earlier, I defined a dynamical system as any system for which we can, in principle, provide a state space evolution equation describing how the values of the state variables of the system change with time. Other values in such an equation specify quantities which affect the behaviour of the system without being affected in turn; these are called the parameters of the system. In formal terms, the coupling of two dynamical systems means that some of the parameters of each system either become, or become functions of, some of the state variables of the other. In the current example, a parameter for the governor is effectively a state variable for the engine, and vice versa (see table i).

TABLE i.

DYNAMICAL COUPLING AND THE WATT GOVERNOR		
	Parameter	State Variable
Watt Governor	Engine Speed	Arm Angle
Steam Engine	Setting of the Throttle Valve: Directly dependent upon arm angle	Engine Speed

Given two coupled dynamical systems, X and Y, as X changes state, the dynamics of Y will change. This change in the dynamics of Y will, in turn, feed back into the dynamics of X, and so on. Effectively, what coupling means is that two systems evolve together through time in a continuous process of feedback and

3 The Internal State of Representationalism

Naturally, the staunch representationalist won't throw in the towel without a fight. In this section, I

and computation which just cannot be applied to the centrifugal governor.

So van Gelder's claim, that the orthodox framework is inappropriate for analysing the interactive subtleties present in the example of the Watt governor, seems to be safe. The crucial step of transferring the same conclusion to the case of cognition rests on the insight that the behaviours of the different devices (governors and cognizers) are equivalent in the sense that both the governing task and cognitively driven human behaviour inherently involve subtle interactions with constantly changing environments. Van Gelder's polemic against the computational analysis of the Watt governor amounts to the charge that traditionalists cannot explain that system's situatedness (the way in which it is embedded in its environment). In this context, the relative simplicity of the governor works in the dynamicist's favour. Assuming that both of the situated behaviours have the same fundamental character, if the traditional style of explanation fails for the Watt governor, it's just bound to collapse in the face of the complexity displayed by cognitive systems and, therefore, will fail to explain their situatedness too.¹⁰

4 The Status of Time in a Cognitive Theory

Van Gelder claims that there is "an important sense in which time does not matter in the operation of the computational governor" (p.8). What does this mean? In this section, I shall endeavour to unpack van Gelder's temporality argument.

Obviously, there is a practical constraint on the computational governor; i.e., it must be successful in its functional niche. Consequently, whatever the necessary sub-tasks turn out to be, they must occur in the right order and happen sufficiently fast. This requires decisions regarding the choice of algorithms and hardware. But, according to van Gelder, these are pragmatic implementation-details, beyond which

...there is nothing which dictates *when* each internal operation takes place, *how long* it takes to carry out, and *how long* the temporal interval between each operation is...The timing of the internal operations is thus essentially *arbitrary* relative to that of any

in the speed of the flywheel caused by some fluctuation in load or pressure. Such observations support the hypothesis that dynamical systems theory is the most appropriate vocabulary in which to analyse the behaviour of the centrifugal governor. In dynamical coupling, remember, the linked systems evolve *together* through time. Hence van Gelder's conclusion that "the temporality of the centrifugal governor is that of the engine itself" (p.9).

Van Gelder is surely right both that cognitive science needs to rethink its attitude towards temporality, and that the likely result will favour the dynamicist. No cognitive agent can be understood without some notion of real-time interaction with an environment, and dynamics, as the science of change, has intrinsic temporality embedded in its theoretical language. This insight can be developed further. Arbitrary temporality is symptomatic of a process of abstraction away from situatedness and embodiment, a process which has been a defining characteristic of most cognitive science (cf. the criticisms levelled at traditional

The first sub-hypothesis requires some clarification. On van Gelder's picture, the Watt governor usurps the Turing Machine as the conceptual anchor for cognitive science. It does so because it fixes the class of non-computational dynamical systems, and *not* because the level of dynamical complexity is equivalent for both governors and cognizers. Van Gelder is perfectly clear on this point (pp.36-37). The dynamical properties of cognitive systems are, of course, far more complex than those of the centrifugal governor. (In section 7 of this paper, I shall explore the nature of cognitive dynamics, arguing, along the way, that the increase in dynamical complexity buys the cognizer a certain sort of independence from her environment. The centrifugal governor does not enjoy such independence.)

The traditionalist is not finished yet. "Look," she says, "I grant you dynamicists that the nervous system *is* (probably) a dynamical system of the sort you describe. If we are talking neurophysiology you have a point. At *that* level of description, you are most likely correct; human beings are dynamical systems. But that doesn't prove to me that *cognition* is best described as state space evolution in dynamical systems."

How might the dynamicist respond? Two possible responses take the form of appeals to current empirical research in cognitive science. However, despite the tactical similarity, the two appeals should be kept apart. The first cites a growing body of empirical

6 The Dynamics of Situatedness

Autonomous agents are fully integrated, self-controlling, active systems which, while in continuous long-term interaction with their environments, behave so as to achieve certain goals. Physical autonomous agents — such as people, animals, insects and autonomous robots — are necessarily embodied, and inextricably situated. Cognitive architectures can be thought of as the control systems for sufficiently complex situated agents (cf. (Harvey *et al.*, 1993)). There now exists a breed of autonomous agent researcher which refuses to take refuge in disembodied, isolated sub-domains of cognition. Instead

the foundations of cognitive science may be in the air. But something has been missing from my discussion. Aside from a few remarks, details of the new path to the cognitive holy grail have been conspicuous by their absence. The traditionalist is owed at least a sketch of the dynamical alternative. It is time to discharge the debt.

7 Where We Go From Here

Along with van Gelder and Beer, I wish to conceptualize an agent and its environment as two coupled dynamical systems. How might we flesh this out? Consider Beer's dynamical framework for autonomous agents (Beer, 1992). A (the agent-system) and E (the environment-system) are coupled via two functions, S (environmental state variables to agent parameters) and M (agent state variables to environmental parameters). As an approximation, S can be thought of as sensory input, and M as motor behaviour, although Beer is clear that the scopes of these two functions are not necessarily restricted to what are normally thought of as sensory and actuation channels. They are supposed to capture *any* effects which one of the systems can have on the future trajectory (state space evolution) of the other. In the rest of this section, I shall investigate the character and consequences of this picture of autonomous agents and their environments, with particular emphasis on those most complex of autonomous agents, cognizers.

Earlier in this paper, I characterized coupling as a relationship in which two systems evolve together through time in a continuous process of feedback and mutual interaction. I propose that dynamical coupling is the fundamental mechanism of situatedness. As a result, agents and their environments become, in some sense, inseparable. In fact, it should be said that they stand in a relation of cospecification (cf. (Varela *et al.*, 1991)). This view is supported by Beer's observation that whilst it is useful, under certain circumstances, to think of the agent and its environment 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 (Beer, 1992). So the decision to draw the line between agent and environment is inherently revisable. Furthermore, if a division is imposed, where the dividing line is drawn is, itself, a matter of choice. To my mind,

¹⁵⁾ Structural instability ensues when a small perturbation results in a qualitative change in the total dynamics of the system. By their very nature, structurally unstable dynamics will tend to be ephemeral. As observers, we may choose to distinguish any number of dynamical cognitive systems and sub-systems at varying levels of abstraction, depending upon the properties (state variables and parameters) we decide to be of current interest. On grounds of overall behavioural coherence, it seems likely that the vast

time slices of the system under investigation (change in parameters = 0), and gradually piece together an overall picture of the way change in the system is changing (Beer, 1992).¹⁸

So I'm prepared to bet that cognitive agents viewed as dynamical systems display a strong tendency towards structural stability, and exhibit highly complex internal dynamics. On the basis of these hypotheses, and accepting Beer's picture of the way in which two coupled systems affect one another's future trajectories, it seems that 'being coupled' does not mean 'being at the mercy of every environmental change'. However, this recognition is perfectly consistent with the view that the behaviour of a cognitive system cannot, in any final sense, be explained in isolation from the dynamics of the environment in which that agent is situated. Behaviour is a feature of a system in which an environmentally-embedded agent and an agent-embedding environment evolve together through time, in a process of mutual and continuous feedback.

It is time to remind ourselves of the charge against which the last two sections of this paper have been a defence. In section 5, our hypothetical traditionalist accused the dynamicist of doing a fine job of theoretical neuroscience, but missing cognition altogether. Some traditionalists might argue that my preferred defence, based, as it is, on a research programme which currently studies relatively simple systems, supports, rather than disposes of, that criticism. This would be justified if my defence of dynamicism rested on the levels of behavioural complexity demonstrated at present by robots (real or simulated) which have been developed by the adaptive behaviour community. But this is not the case. Here's a reconstructed summary of my argument:

1. Real cognitive agents are necessarily both situated and embodied. (In the case of animals, 'embodiment' is understood as referring to sensorimotor capacities.)
2. Cognitive architectures are the control systems for the situated activity of sufficiently complex embodied agents.
3. Thus, cognition cannot be studied effectively in isolation from situatedness and embodiment. (Or-

framework is ultimately a matter for empirical research to decide, and Beer is often rather measured in his attacks on traditionalism. But I, for one, bet in favour of the view that traditional styles of explanation will become increasingly marginalized as the field of adaptive behaviour reaches maturity.

8 Summary

In this paper, I have defended the emerging dynamical approach to cognitive science. The space of possible dynamical systems is immense, and representational/computational systems fill one tiny corner of that space. There is mounting evidence that explanations couched in the vocabulary of representations and computations are far too restrictive to account for the behaviour of cognitive systems. The most compelling arguments for a more general, dynamical perspective come from the study of the adaptive behaviour of situated, autonomous agents. Once situatedness and embodiment are firmly on the agenda (as opposed to being tucked away as the sort of ‘any other business’ which always gets left until the next meeting), dynamical systems theory looks to be the most promising theoretical language in which to conduct cognitive science.

Acknowledgements

This work was supported by British Academy award no. 92/1693. Many thanks to Maggie Boden, Dave Cliff, Matthew Elton, Inman Harvey, Phil Husbands, Geoff Miller, and Chris Thornton for their helpful comments on an earlier version of this paper.

References

- Abraham, R. H., & Shaw, C. D. 1992. *Dynamics - The Geometry of Behaviour 2nd edition*. Redwood City, California: Addison-Wesley.
- Baker, G. J., & Gollub, J. P. 1990. *Chaotic Dynamics - An Introduction*. Cambridge: Cambridge University Press.
- Beer, R. 1990. *Intelligence as Adaptive Behaviour: An Experiment in Computational Neuroethology*. San Diego, California: Academic Press.
- Beer, R. 1992. *A Dynamical Systems Perspective on Autonomous Agents*. Tech. rept. 92-11. Case Western Reserve University, Cleveland, Ohio.
- Brooks, R. 1991a. Intelligence Without Reason. *Pages 569–95 of: Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*. San Mateo, California: Morgan Kaufman.
- Brooks, R. 1991b. Intelligence Without Representation. *Artificial Intelligence*, **47**, 139–59.
- Brooks, R. 1992. Artificial Life and Real Robots. *Pages 3–10 of: Varela, F. J., & Bourgine, P. (eds), Toward a Practice of Autonomous Agents: Proceedings of the first European Conference on Artificial Life*. Cambridge, Massachusetts: MIT Press / A Bradford Book.
- Cliff, D. 1990 (May). *Computational Neuroethology: A Provisional Manifesto*. Cognitive Science Research Paper CSRP162. University Of Sussex.
- Cliff, D., Husbands, P., & Harvey, I. 1993. Analysis of Evolved Sensory Motor Controllers. *Pages 192–204 of: Proceedings of the 2nd European Conference on Artificial Life*.
- Dawkins, R. 1982. *The Extended Phenotype*. Oxford: Oxford University Press.

- Gibson, J. J. 1979. *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.
- Giunti, M. 1991. *Computers, Dynamical Systems, Phenomena , and the Mind*. Ph.D. thesis, Department of History and Philosophy of Science, Indiana University.
- Goldberg, D. E. 1989. *Genetic Algorithms in Search, Optimization and Machine Learning*. Reading, Massachusetts: Addison-Wesley.
- Harvey, I., Husbands, P., & Cliff, D. 1993. Issues in Evolutionary Robotics. *Pages 364–73 of: Meyer, J. A., Roitblat, H. L., & Wilson, S. (eds), From Animals to Animats 2: Proc. second International Conference on simulation of Adaptive Behaviour*. Cambridge Massachusetts and London, England: MIT Press / A Bradford Book.
- Heidegger, M. 1926. *Sein Und Zeit (English translation by J. Macquarrie and E. Robinson. 1962. Being and Time)*. Oxford, England: Basil Blackwell.
- Husbands, P., Harvey, I., & Cliff, D. 1993. *Circle In The Round: State Space Attractors for Evolved Sighted Robots*. Unpublished Manuscript ,University of Sussex.
- Maes, P. (ed). 1990. *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*. Cambridge, Massachusetts: MIT Press.
- McAuley, J. D., Anderson, S., & Port, R. F. 1992. Sensory Discrimination in a Short-Term Dynamic Memory. *In: Proceedings of the 1992 Cognitive Science Society*.
- Simon, H. A. 1969. *The Sciences of the Artificial*. Cambridge, Massachusetts: MIT Press.
- Simons, P. 1992. *The Action Plant*. Cambridge, Massachusetts, and Oxford, England: Blackwell Books.
- Smithers, T. 1992. Taking Eliminative Materialism Seriously: a methodology for autonomous systems research. *Pages 31–40 of: Varela, F. J., & Bourgnine, P. (eds), Toward a Practice of Autonomous Agents: Proc. 1st European Conference on Artificial Life*. Cambridge, Massachusetts:

van Gelder, T. 1992. *What Might Cognition be if not Computation?* Tech. rept. 75. Indiana University Cognitive Sciences.

Varela, F. J., Thompson, E., & Rosch, E. 1991. *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, Massachusetts, and London, England: MIT Press.