

Does Embeddedness Tell Against Computationalism? A Tale of Bees and Sea Hares

Matteo Colombo^{†§}

Abstract There is a tendency in the cognitive sciences to emphasize that cognition is embedded in a world and intrinsically embodied. In a sense it is uncontroversial that cognition is embedded. However, it is not obvious what the best explanatory framework to understand embeddedness is. Tim Van Gelder and Randy Beer, among others, take it that embeddedness calls for a dynamics perspective and at the same time tells against computationalism: If we want to understand how brains are coupled with their environments, the dynamicist framework is the best, since the traditional computational one seems to be inadequate to capture this kind of phenomena. This paper gives grounds to reject this argument. By focussing on two case-studies, it argues that, on the one hand, the dynamicist framework is not sufficient to understand embeddedness; on the other, the computationalist framework is resilient and is still necessary even to understand a phenomenon involved in embeddedness such as brain-environment coupling. The conclusion is drawn that embeddedness does not provide telling evidence against computationalism.

1 INTRODUCTION

Presumably, no one would question that cognition is, in a sense, situated (or embedded).¹ The environment where an agent is situated plays an important role in its cognitive activity and behaviour. It is also uncontroversial that the relation between an agent and its immediate environment is of one of ongoing interaction. When I say that a cognitive capacity is situated, or embedded, I intend that the processes that underlie that capacity take place and develop when a coupled system emerges from the complex, real-time interplay between brains, environment (and bodies). The adaptive success of a situated agent depends on the kind of causal coupling between its brain, and environment (and body).

Which explanatory framework is best to understand embeddedness is controversial however. Different approaches to cognitive sciences give different force to the claim that cognition is embedded. Tim Van Gelder [17] for example considers embeddedness as an argument in favor of a dynamical approach

to cognitive science and *at the same time* a recalcitrant case for the computational approach. "Dynamical cognition – he claims ([17], p. 623) – sits comfortably in a dynamical world." From this claim alone, however, it doesn't follow that the dynamical framework is the best to understand embeddedness. The aim of this paper is to provide some ground for why a dynamical framework may not be sufficient for *understanding* embeddedness. A computational framework still yields necessary insight to understanding how brains and world interact.

Here's how I plan to proceed. Section 2 delimits the playground of my argument. It clarifies the terminology, and makes it explicit the assumption of the argument. Section 3 recalls two case-studies on associative learning. I take it that learning is a paramount example of cognitive ability where embeddedness matters for the following reason. Learning, broadly understood, depends on the interaction between an agent's brain (and body) and its surroundings. By "observing" its interaction with the environment with which its brain is coupled, an agent gains information that enables it to improve its future decisions. Because it increases the "appropriateness" of a behavioural response to a given class of environments, learning contributes to adaptive behaviour. Strictly speaking, then, the target of my argument is whether dynamicists or computationalists give the best explanation of one example of embedded cognitive capacity, namely learning.²

Section 4 elaborates on the case-studies by arguing for the necessity of a computationalist framework (and against the sufficiency of the dynamicist one) to understanding embeddedness. Section 5 concludes the paper by summarizing results so far, and drawing the moral that those who think that a dynamicist framework is sufficient for understanding embeddedness may be mistaken.

2 DYNAMICISM, COMPUTATIONALISM AND EMBEDDEDNESS

Because there are various ways in which different theorists conceive of the key terms, it is too easy to be trapped in terminological misunderstandings and lose sight of the

[†] Dpt. of Philosophy, The University of Edinburgh, UK.

[§] CRESA (Center for Experimental and Applied Epistemology), Milan, Italy

Email: M.Colombo-2@sms.ed.ac.uk

¹ Embeddedness is often taken to be one of the hypotheses belonging to a family of approaches to understanding cognition known as "Situated cognition movement" [19]. Since in this paper nothing hinges on the distinction between situatedness and embeddedness, I will use 'situatedness' and 'embeddedness' interchangeably.

² The assumption that learning is a cognitive capacity that qualifies as embedded is not obvious however. Learning, in fact, covers a diversity of cognitive capacities supported by a variety of mechanisms. It may be that in some cases of learning the role played by the environment is not crucial. And therefore the issue of embeddedness would be something as a red-herring in those cases. Although I shall not put forth a general argument for why, and in what cases, learning qualifies as embedded, I shall motivate why the cases exposed in section 3 can be regarded as cases where embeddedness, or features that characterize embedded cognition (e.g. coupling) are in fact important.

substantial issues in the “dynamicism - computationalism debate”. It is therefore necessary to make some terminological clarifications. This section is devoted to put some bits of terminology in place.

Dynamical system theory is a mathematical theory that provides us with analytical tools to study the behaviour of complex systems. Taking a dynamicist perspective on a certain *explanandum* means to adopt specific concepts, metaphors, a vocabulary to frame our understanding of that phenomenon. The core concepts that form the dynamicist framework are the ideas of state space, time set, and evolution operator. These concepts enable us to draw a *geometrical* analysis of the phenomenon we seek to explain. Let a system be a set of interdependent variables that define an *explanandum*, the state of the system at a time is the value of its variables at that time. The state space of the system consists of the possible overall states of the system through time. The state space can be of any topology and dimension according to the number and nature of the variables of the system. The behaviour of the system is understood in terms of the trajectory of the system within the state space through time. The trajectory is governed by the evolution operator which is typically a set of differential equations. Concepts from dynamical systems theory are becoming increasingly important in cognitive science [2]. Dynamical system theory tools are for example among the nuts and bolts of computational cognitive neuroscience [12]. Nevertheless, this does *not* mean that the dynamical *framework* has replaced, or is even near to replace, the “old-fashioned” computational one.

The theory of computation is also a body of mathematics. The adoption of the analytical tools of the theory of computation to study cognition and behaviour of an agent is said ‘computationalism’. More precisely, computationalism holds that the cognitive processes and behaviour of an agent is explained by computations. A computation is the manipulation of symbols according to some algorithm. An algorithm is a step-by-step procedure for accomplishing something. A symbol is a carrier of information, it represents something. A related, but conceptually distinct, broader, construal views computation as information-processing. Accordingly, computation is what is involved in manipulating, storing, retrieving, encoding, trafficking information – whatever it might be. If we apply this framework to cognition, cognition and behaviour of an agent are understood in terms of patterns of information transformed, retrieved, stored, processed by a mechanism according to some input-output function. This is what I mean with “computational framework”.

A number of disclaimers are in order at this point. First, as mathematical theories the dynamical and the computational theory are roughly equipotent [3]. Therefore, the claim that dynamicism is a better framework than computationalism for understanding embeddedness is not supported by mathematical reasons alone. Second, the characterizations above are “cartoon” characterizations in two senses. On the one hand, they overlook distinctions that motivate different positions within the same approach. For example, Van Gelder’s dynamical hypothesis [17] has raised criticism not only from the computationalist party, but also within the dynamicist group. On the other hand, both dynamicism and computationalism face theoretical and methodological “internal” challenges (e.g. [11]; [2], p.114-117). For my purpose however, suffice the cartoon characterization. Third, I assume that *connectionism* represents a refinement of

computationalism *and* is in continuity with dynamicism.³ On the one hand, in connectionism, computations can be seen as distributed across neural networks. On the other, connectionist models can implement equations from dynamical system theory. Finally, and importantly, I take it for granted that cognition is always embedded in an environment. That is, I assume that biological nervous systems are always coupled with an uncertain and changing environment: The brain and the world are causally coupled, and both contribute (with the body) to intelligent behaviour, and cognition. Notice that these are *ontological* claims. Which *explanatory perspective* is the best to understanding embeddedness is a separate issue.

I am not interested in the *nature* of the contribution of the world to cognition. I am interested in assessing whether embeddedness may constitute a really telling argument for assuming a certain *perspective* in the cognitive sciences. A really telling argument for the claim that perspective *X* is better than *Y* (in our case, that dynamicism is better than computationalism) is an argument that supports perspective *X*, and that *at the same time* tells against the rival perspective. If it turns out that computationalism misrepresents embeddedness, whereas the dynamicist framework renders an insightful image of the same phenomenon, then we have an excellent reason to withdraw the computational framework in favor of the dynamicist for understanding embeddedness.

With ‘perspective’, or ‘world-view’ - as adopted by Beer [1], I mean a conceptual framework to understanding a phenomenon. A framework comprises a vocabulary, a set of concepts, metaphors, insights that enable us to understand a certain phenomenon. Thus, the computational framework draws on the metaphor of the brain as a computer, and emphasizes the role of functional information-processing structures. Its vocabulary comprises concepts like “manipulation”, “processing”, representation”, “retrieval”, “storing”, “activation”, “input-output function”. The dynamicist framework draws on the metaphor of cognition as movement. Its vocabulary comprises concepts like “attractors”, “transients”, “coupling”, “bifurcation”, “emergence”, “state space”, “trajectory”.

We should be careful in spelling out what the opposition between dynamicism and computationalism amounts to here. As mentioned above, the opposition is not mathematical, and has not to do with mutual inconsistency. The opposition, instead, is about the different kind of “explanatory priorities”, or of “explanatory concerns” suggested by the two frameworks [4], [6], [7]. Taking a computationalist perspective to understanding some system means to focus on the function that the system computes, to try and identify how certain states of the system stands-in, or represents some states of affairs, what class of algorithms transforms these representations, what constraints there are on the information-processing of the system. A dynamical perspective, in contrast, would set a very different class of priorities in one’s explanatory agenda. A dynamicist would try to identify the relevant set of variables, and parameters that define the state space of the system. He would be concerned both with the evolution of the system in that state space and with the class of equations that can account for the spatiotemporal

³ It would be interesting to consider whether connectionist framework is “the best of two worlds” as explanatory framework. Despite its intuitive appeal, the exploration of this possibility goes beyond the scope of this paper.

trajectory of the system. He will be interested in attractors, repellers, phase portraits.

Obviously, *ceteris paribus*, the unfamiliarity of a framework is not a sufficient reason to reject it for another. We have to examine the different kinds of *understanding* provided by the perspectives under scrutiny. To make the point concrete, compare Aristotelian physics with Galilean. In the XVII century, the Aristotelian framework was more familiar than the Galilean one. Its vocabulary comprised “fire”, “air”, “water”, “earth” (four terrestrial elements), and “circular”, “up” and “down” (the differential motional natures of the elements). The Galilean framework comprised only one element, “corporeal matter”, and different parameters to describe its properties and motions. Facing the same phenomenon, Aristotle talked of a swinging stone striving to reach its natural resting place; instead, Galileo talked of a pendulum, a periodically moving body, whose movement can be understood in terms of frequency, amplitude, radius of the pendulum. The difference is not in precision. The difference is conceptual: By adopting a different vocabulary, the Galilean framework provided better explanatory understanding on the physics of pendula. The ultimate reason why it superseded the Aristotelian framework is empirical: Empirical success always leads the way in the choice of one framework over another.

3 DYNAMICISM AND COMPUTATIONALISM AT WORK

Embeddedness bears on a fundamental cognitive ability: learning. This section centers on two case studies where adaptive behaviour of an agent arises from its ongoing interaction with its environment. Both cases are from the field of computational neuroscience,⁴ and both are concerned with learning abilities of simple organisms. The first implicitly assumes that brains are kinds of computers, thereby adopting a computationalist perspective; the second frames its results in dynamicist terms. The two cases in turn.

3.1 HONEYBEES

How may honeybees learn what flowers to visit for getting their next meal? Montague and colleagues [13] tackle this problem by constructing a model of a bee foraging in an uncertain environment. They draw on behavioural observations and neurophysiological data. It seems that bees' foraging behavioural repertoire is based on associations between the occurrence of stimuli (e.g. color of flowers) and outcomes (amount of nectar yielded by the flower). Through trial and error interactions, bees establish proper associations: The more nectar, the more likely bees will return to that flower. Here the key notions to understanding bees' foraging behaviour are *prediction* and *reward*. A bee anticipates what its internal state and the external world will be like by using its current and past experience of reward. *Rewards* are used to improve the quality of predictions. ‘Reward’ can be defined operationally as the positive value that a system places on the attainment of a certain

goal. This kind of learning seems to be supported by a neuron in the bee ganglion, the VUMmx1, which releases octopamine. The activity of this neuron seems in fact to encode a prediction of reward which enables the bee to improve its performances.

Montague and colleagues' model simulates the behaviour of a bee foraging over a virtual field of flowers. The bee is endowed with a visual system that processes inputs from the environment and represents changes in percentages of color. The goal of the bee is to get nectar. To facilitate the attainment of this goal, its computational system guides the bee over the field. The computation is based on current sensory inputs and a prediction of nectar-reward built on remembered rewards associated with different states of the environment, that is, with different colors. The activity of VUMmx1 enables the bee to predict reward by computing prediction-errors with its activity. After the bee has landed on a flower, the neuron combines information about the current reward (the amount of nectar yielded by that flower) with its own prediction (what reward it expected from that flower). Then, it transmits information about how well the actual reward tallies with the predicted one to the rest of the system. When the actual reward is better than expected, VUMmx1 output leads the rest of the system to upgrade the value attached to the state that yielded that amount of reward. In other words, it gives a motive to the bee to remember that the color of that flower predicted an amount of nectar better than expected, and to use this courtesy of the prediction-error system, the bee learns to choose the most adaptive actions by interacting with its environment. Let's turn to another case now.

3.2 SEA HARES

How may sea hares learn to bite edible food and avoid inedible food? Phattanasri and colleagues [15] focus on this food-edibility problem drawing on previous observations of the behaviour of *aplysia*, a genus of sea hares. They model sea hares as agents equipped with a mouth, a smell sensor, and a gut sensor. The goal of this agent is to learn to eat only edible food in a changing environment containing either edible or inedible food. To reach this goal the agent has to learn to associate the right smell to the right type of substance and take an action accordingly by relying on its experiences in that environment.

Phattanasri and colleagues show that a continuous-time recurrent 3-neuron network lacking plastic synapse can evolve to solve this task. This kind of agent is not endowed with a specific learning mechanism; it evolves its learning ability by using the stream of binary smell-sensory inputs from its environment, and the gut sensor serving as a reinforcement signal. Thereby, the evolution of adaptive smell-sensitive actions is function of reinforcement construed as a dynamic property of the [environment-agent-food type] system. Accordingly, they give a topological analysis of the evolution of the system. First, for each of the five possible input patterns (i.e. “no input”, “good smell”, “bad smell”, “positive reinforcement”, “negative reinforcement”), the complete phase portrait of the circuit is determined.

The understanding yielded on the learning task under consideration is in fact in terms of phase portraits, that is, of a plot of trajectories in the state space of the 3-neuron circuit. In this way, attractors, basin, and stable equilibrium points are revealed. Then, as the input signal varies over time, the circuit state is observed to move through different phase portraits attracted towards the equilibrium points identified beforehand.

⁴ Broadly, computational neuroscience is the use of mathematical modeling, and computer simulations to understand the brain. By itself, this does not entail a commitment to computationalism. However, many computational neuroscientists do make the assumption that brains are kinds of computers (e.g. [5]; [9]).

The dynamic of the system through phase states generates changes in behaviour such that, even though the agent is not endowed with a specific learning mechanism, it can successfully learn to eat only edible food.

4 A PLEA FOR COMPUTATIONALISM

The two cases above serve as bedrock where to build my case for computationalism. My argument is in two parts. First, I examine one reason in favor of the sufficiency of the dynamicist framework, and I rebut it. Then, I motivate why the computationalist framework is still necessary to understand embeddedness.

4.1 ON COUPLING

Let's assume that brains and their surrounding environment are coupled. Coupling is a continuous reciprocal, causal dependence: That a system is coupled with its surroundings means that the system both affects and is affected by what surrounds it [6], [17].

Coupling is usually taken as a reason in support of the arbitrariness of distinguishing brain-centered cognitive systems from the environment where they are embedded [3], [17]. To understand the interactive complexity underlying embeddedness, so runs the argument, we should adopt a dynamicist perspective. In fact, when it comes to *understanding*, relating brain and environment by conceiving of them as a single system is less problematic than relating systems of different kinds. According to this argument, the learning ability of the sea hare is best understood as a dynamics of the system [food-smell sensors-gut sensors] evolving towards an adaptive equilibrium. The general moral is that learning may not be either a behavioural or neural natural kind, but rather, a systemic ability [15]. But does coupling give a strong reason to embrace a dynamicist framework? Or, put differently, does a dynamicist perspective suffice to understand the complex interplay between brains and environment? There are two reasons why it is problematic to affirmatively answer these questions. The first has to do with "componential analysis", the second with "comparative understanding".

First, it may be difficult to understand in dynamical terms the specific, partial contributions of components of the systems to the evolution of the learning ability.⁵ The dynamicist may tell us that the contribution of the smell sensors of the sea hare consists in the influence that their values have on the phase portraits in the space state of the system. But this is unsatisfactory: We would like to understand the functional,

information-processing role of that component during the evolution of learning. In Montague and colleagues' simulation, we know that learning develops in virtue of an internal supervisor that assesses the ongoing performance of the bee in light of its goals. The bee "is teaching itself about its world" from the feedback of its actions ([16], p.104). Therefore, in accounting for an embedded capacity as learning a dynamical perspective may be insufficient since it would provide little understanding in the information-processing machinery that supports the evolution of the learning of the sea hare.

At this point, the dynamicist may have two objections. He may first point out that it is not obvious how the case of associative learning in the bee is an example of situated cognition where coupling matters. In fact, the bee with its actions doesn't really affect the environment where it is embedded. Why then should we consider this as a case of embedded cognition? In the second place the dynamicist may object that we are simply begging the question: In making this request, we are assuming that the system has to be understood in computational terms. But, according to the dynamicist, it is arbitrary to understanding learning as specifically linked to the information-processing role of one component, rather than to the dynamics of the [brain-body] system.

The response to the first objection goes as follows. The learning of the bee can be considered as an example of situated cognition because the bee does not passively retrieve perceptual information from its world. The representation of the environment that the bee has constantly changes as it makes decisions by drawing upon the values attached to the representations. Thus, certain environmental features (e.g. different patches of color) are really re-constructed depending upon the goal-oriented actions of the bee. By interacting with its world, the bee actively constructs a "value-laden" representation of its environment, which in this sense can be taken to be affected by the activity of the bee. What is the coupled system then? The obvious place to look for causal loops in the bee case is the causal relationship between its value-laden representations of external states of affairs and the decision it makes. The bee's perception, that is, the bee's representation of the external world, affects its decisions which in turn affect perceptions, and so on. Ultimately this kind of complex interplay driven by the prediction-error system leads the bee to display adaptive behaviour.

The second objection can be answered by pointing out to the dynamicist that we have good, *independent* reason to ask about information-processing roles. We regard certain systems as coupled precisely because of that role. A mechanism as gut-sensor is *taken* to be coupled to the sea hare environment because we have a computational pre-understanding of its role. Without having this kind of understanding it would be problematic to identify where to apply the dynamics analysis – whether at the level of brain-body-environment system, or of body-neuromechanical interactions, or neural interactions. For we wouldn't have an *independent* rationale to understanding why we should (de)couple the system in certain ways rather than others. For example, it may be not a good idea to put forth a dynamical analysis of the cognitive ability involved in the conversation you carry on at a crowded pub by focusing on the coupled dynamical system [brain-pub environment]. For in this case it may be more revealing to unfold the computational machinery employed by your brain to pull the right signal out of

⁵ Strictly speaking, the sea hare agent modeled by [15] is *non-autonomous*. An autonomous system produces its control signals without benefit of external sensory inputs. Since it receives time-varying inputs, the sea hare agent is not autonomous. However, when it comes to understanding the behaviour of the agent, Phattanasri and colleagues consider it as an autonomous dynamical system by analyzing the sea hare-environment dynamics holding the input fixed to certain values. The coupling here is taken to be a useful *epistemic device* to best understanding the evolution of the agent. One may then wonder whether this study makes a really strong case for dynamicism. Although, perhaps, the sea hare case is not the strongest one for supporting a dynamicist approach to embeddedness, it fits well with my overall argument. In fact, my focus is precisely on *explanatory framework*, on what *epistemic device* is best to understand certain embedded cognitive abilities.

the buzz all around you. For these reasons, it is problematic to take it that coupling is a strong reason to favor a dynamicist approach.

The second problem for a dynamicist framework on embeddedness has to do with the understanding of cognitive analogies across different kinds of agents. Why do agents embedded in different environments (and embodied in very different bodies) seem to display analogous cognitive abilities? Consider the ability of the honeybees to make reward-based predictions and act on its basis. This is an adaptive cognitive strategy that seems to be displayed also by monkeys, rats, and humans. A dynamical framework may be insufficient to understanding this kind of analogy. For, if we treat the agent and its environment as a single complex coupled system, then a cognitive ability has to be understood also in function of the environmental (and bodily) variables of the coupled system. These details are different for bees, rats, humans. Hence, for each kind of agent, we would have different stories about what seems to be the same ability. If this is so, then understanding an apparent analogy would be problematic. In this case, to understand the interplay between brains and environment a unifying, computational framework is best. In fact, a wealth of behavioural and neural data suggests that reward-prediction learning is an ability supported by a particular computation across different kinds of agents. Both in humans, rats, and honeybees, the firing of dopamine (or in bees, a similar chemical called “octopamine”) cells seems to encode prediction-error signals governed by a temporal prediction algorithm that accomplishes a specific computational task [14]. A dynamical framework, therefore, may be insufficient to compare abilities across different agents, and thereby grasping analogous cognitive strategies underlying the behaviour of different agents.

4.2 DOING WITHOUT COMPUTATIONS?

The considerations above give some suggestions for why a computational framework is still necessary to understand embeddedness. Let’s assume that a computational talk is not only insufficient, but also is unnecessary for an account of embeddedness. We must not talk of the computation of input–output functions in the honeybee case. We must refrain from construing the functions being computed in terms of representations, that is, in terms of the information content carried by the electric signals travelling on a neural network – for instance, information about food smells, or “appropriate” behaviour in a certain kind of environment. If an array of cells computes a prediction-error, we must not call it computation.

It would be unclear what we would gain by framing our understanding of embeddedness purely in dynamical terms. Certainly we would lose grip on the functional role of the components of the system.⁶ Dispensing with computational talk may be problematic for a reason of “expressive convenience” as well. For example, in *Science Without Numbers* Hartry Field shows how it is possible to do Newtonian physics without

⁶ Another way to make this point might be in terms of Dennett’s distinction between the physical stance and the design stance. The dynamical framework fails to capture certain properties of systems that become discernible once we offer computational explanations of the same behaviour. This might be akin to Dennett’s point that certain explanations become available once we adopt the design stance that are not available from the physical stance [10]. Thanks to Julian Kiverstein for pointing this out to me.

mathematical statements; however, it would be very inconvenient for a scientist, to say the least, if she did without mathematical statements. This possibility is not a sufficient ground to do without numbers.

It might be objected that in all these cases we are unfairly focusing on decoupled brains thereby suggesting a privileged cognitive role for brain-based cognition. This objection, however, misses the target. On the one hand, if we want to understand how brains relate to the environment where they are embedded, we have to understand how the brain component of the brain-body-world system works. And to understand how brains work, the computational framework still seems necessary. On the other hand, computationalism by itself doesn’t entail an individualistic “brain-bound” view of cognition. As argued by Wilson [18] and by Clark & Chalmers [8], computational systems that support cognition can extend beyond the skull. The point made above bears on the necessity of a computational framework; a separate issue is whether this framework is better applied narrowly, to understanding the brain alone, rather than widely, to understanding extended cognitive systems of which brains are components. My argument is not meant to have bearing on this issue.

5 CONCLUSION

The Dynamical System Theory is already part of the toolbox of cognitive science. And for good reason since it is an excellent analytical tool to deal with complexity and to have a geometrical analysis of it. However, it is not obvious that dynamics provides the best *conceptual framework* with which to understand cognitive processes. Some (e.g. [1], [4], [15], [17]) suggest that dynamicism is preferable for the study of embedded agents. Embeddedness – the argument runs – provides reason to prefer dynamics over computationalism as conceptual framework. This paper has tried to challenge this argument by examining two cases where the learning of a simple agent interacting with an uncertain environment is understood within different frameworks.

The main claims made and defended so far are five.

- 1) Even if cognitive systems *are* dynamical systems, it is not obvious that cognitive systems are best *understood* in dynamical terms.
- 2) Learning is, in a sense, a paradigmatic case of embeddedness.
- 3) From 2) it doesn’t follow that at least certain cases of learning are best understood in dynamical terms.
- 4) Embeddedness involves coupling. But coupling is not a sufficient reason to prefer dynamics over computationalism.
- 5) Computationalism is still necessary to understanding (at least certain aspects) of embeddedness.

The conclusion follows that embeddedness may *not* tell at the same time against computationalism and for dynamics. As always in science, empirical results will tell the last word on this issue. If as a conceptual framework computationalism systematically prejudices, thereby biasing, the answers to empirical questions about embedded cognition, then we will have an excellent reason to prefer an alternative framework.

ACKNOWLEDGEMENTS

A sincere thank you to Andy Clark and Julian Kiverstein who provided useful comments on a previous draft of this paper, and to Danielle Brown, who assisted me with the editing of the essay. Obviously, the usual disclaimers about the remaining errors in the paper apply.

REFERENCES

- [1] Beer, R.D. (1995). "Computational and dynamical languages for autonomous agents". In R. Port and T. van Gelder (Eds.) (pp. 121-147).
- [2] Beer, R.D. (2000). "Dynamical approaches to cognitive science". *Trends Cog. Sci.* 4:91-99.
- [3] Beer, R.D. (2008). *The dynamics of brain-body-environment systems: A status report*. In P. Calvo and A. Gomila (Eds.), *Handbook of Cognitive Science: An Embodied Approach* (pp. 99-120). Elsevier.
- [4] Beer, R.D. (in press). *Dynamical systems and embedded cognition*. To appear in K. Frankish and W. Ramsey (Eds.), *The Cambridge Handbook of Artificial Intelligence*. Cambridge: Cambridge University Press.
- [5] Churchland, P. S., & Sejnowski, T. J. (1992). *The Computational brain*. Cambridge, MA: MIT Press.
- [6] Clark, A. (1997). *Being There: Putting Brain, Body and World Together Again*. MIT Press.
- [7] Clark, A. (1997). "The Dynamical Challenge" *Cognitive Science* 21:4:1997 p 461-481
- [8] Clark, A., & Chalmers, D. (1998). "The Extended Mind". *Analysis* 58:10-23.
- [9] Dayan, P (1994). "Computational modelling". *Current Opinion in Neurobiology*, 4:212-217.
- [10] Dennett, D. (1987). *The Intentional Stance*. The MIT Press.
- [11] Floridi, L. (ed.) (2004). *The Blackwell Guide to the Philosophy of Computing and Information*. Oxford - New York: Blackwell.
- [12] Izhikevich, E.M. (2007). *Dynamical Systems in Neuroscience*. Cambridge, MA: MIT Press.
- [13] Montague, PR, Dayan, P, Person, C, Sejnowski, TJ (1995). "Bee foraging in uncertain environments using predictive Hebbian learning". *Nature* 377:725-728.
- [14] Niv, Y., & Montague, R.P. (2009). "Theoretical and empirical studies of learning". In: PW Glimcher et al, (Eds.), *Neuroeconomics: Decision making and the brain*, Chapter 22, 329-249, Elsevier.
- [15] Phattanasri, P., Chiel, H. J., & Beer, R.D. (2007). "The dynamics of associative learning in evolved model circuits". *Adaptive Behavior*, 15: 377-396.
- [16] Quartz, S.R., & Sejnowski, T.J. (2002). *Liars, Lovers, and Heroes: What the New Brain Science Reveals About How We Become Who We Are*. New York: Harper Collins Publishers Inc.
- [17] Van Gelder, T. (1998). The dynamical hypothesis is cognitive science. *Behva. Brain Sci.* 21:615- 628.
- [18] Wilson, R.A., (1994). "Wide Computationalism". *Mind* 101: 351-372.
- [19] Wilson, R.A., & Clark, A. (2009). "How to Situate Cognition: Letting Nature Take its Course," in P. Robbins and M. Aydede (eds.) *Cambridge Handbook of Situated Cognition*, Cambridge: Cambridge University Press, Ch. 4.