No Cognition without Representation? Dynamical Computationalism and the Emulation Theory of Representation

Samir Chopra
Department of Philosophy
Graduate Center of the City University of New York
25 West 43rd Street
New York, NY 10036
schopra@csp.gc.cuny.edu

May 13, 1999

Abstract

The dynamical hypothesis threatens computationalism and its commitment to mental representations: the continuous, reciprocally causative nature of cognition renders impossible such an analysis. Andy Clark and Rick Grush have suggested a hybrid approach that models cognitive systems using dynamical systems theory but employs mental representation via *emulation* to aid its explanations. I will call such a view *dynamical computationalism* and in this paper, I evaluate its prospects.

The proposed 'architecture' of representation, which a cognitive agent *must* possess in order to be distinguished from a merely adaptive one, uses a notion of *partial* agent-environment coupling to facilitate its use of representations. In the *capacity* for emulating the environment for use as internal representation, dynamical computationalism claims to have laid down a *necessary* condition for cognition.

I will show that this approach reveals a misunderstanding of the force of the dynamical hypothesis. I argue too, against the plausibility

and the necessity of the condition laid down: it is framed in terms of a capacity for emulation, but its absence is compatible with cognitive behavior. Dynamical computationalism would be better off without such an implausible condition on cognition elevated to the status of a necessity.

1 Introduction

The dynamical hypothesis' (Thelen and Smith, 1994; van Gelder, 1995; van Gelder and Port, 1995) (henceforth, *DH*) insistence on the temporal embeddedness of cognition and the tight coupling of the cognitive agent with its environment threatens computationalism and its commitment to mental representations: the continuous, reciprocally causative nature of cognition renders impossible such an analysis. It has been suggested that it's possible to reconcile the *DH* and computationalism by[1]:

> ..a general framework that accepts many of the basic premises of..dynamical systems theory

and by[2]:

> ..a hybrid approach that both tracks the complex interactive dynamics and makes some attempts to reveal the information processing roles of component states and processes.

I will call such a view dynamical computationalism (henceforth *DC*) (Clark, 1997, 1998; Grush, 1997a, 1997b; Clark and Grush, 1999)[3]. It contends that the anti-representationalism of the *DH* renders it *incomplete* with respect to *explanations* of cognitive systems: its modeling of cognitive systems works purely at a descriptive level. The best explanation, it is

[1](Clark and Grush, 1999)

[2](Clark, 1998), page 374

[3]Grush (1997a) terms this thesis the Emulation Theory of Representation while Clark and Grush (1999) term it Minimal Robust Representationalism. For the sake of brevity, I will refer to the two as just *DC* throughout. This captures the idea of an approach where all the modeling is done using dynamical systems theory but where representations make an appearance to make the overall approach computational.

claimed, of cognitive systems, whose actions are best understood as information processing, or representation transforming, is provided by a computationalism with an enlightened version of representation.

DC's 'architecture' of representation, which a cognitive agent *must* possess in order to be distinguished from a merely adaptive one, uses a notion of *partial* agent-environment coupling to facilitate its use of representation via *emulation*. In the capacity for emulating the environment for use as internal representation, *DC* claims to have laid down a *necessary* condition for cognition.

I will show that the *DC* approach reveals a misunderstanding of the force of the argument made by the *DH*: in particular, the arguments *against* the existence of arbitrary boundaries between cognitive controllers and representations and *for* the presence of tight coupling in *all* cognitive interactions. I argue against the plausibility and the necesssity of the condition that *DC* lays down on cognition: it is framed in terms of a *capacity* for emulation, but its absence is compatible with cognitive behavior. The *DC* theorist, committed to understanding cognition via a *primarily* dynamical model, would be better off without such an implausible condition on cognition elevated to the status of a necessity.

This is not to suggest that the *DH* has to make no concessions; it does. It has to explain how distinctions are made between kinds of adaptive systems– those that are merely adaptive, those that are somewhat intentional and those that are clearly cognitive. However, this is a problem that the dynamical hypothesis does not *claim* to have solved as yet; its non-triviality is acknowledged by proponents of the *DH*.

2 The Challenge to Computationalism

Computationalism (with respect to human cognition) is the thesis that the brain is formally equivalent to a Turing machine. It treats the transformations of mental representations (as sentences, sentence constituents or patterns of activations) via rules of inference as computation. It bears a commitment to a *discrete* modeling of mental states: the real or virtual machine is decomposed into a set of discrete component states and processes. These states and structures are interpreted in a way that represents causal cycles as being informationally guided and motivated. The processing of information bearing inner states regulates transitions between cognitive states[4].

The dynamical hypothesis insists that cognitive systems are best understood and modelled by using the resources of dynamical systems theory[5] (henceforth, *DST*), and is committed to the following:

1) The continuous, in real-time, nature of cognitive processes;
2) The interdependent, *tightly coupled*, highly complex, mutually causal interactions in which cognitive agent and environment participate.

The first is clearly antithetical to computationalism. The second creates the continuous reciprocal causation (henceforth, *CRC*) problem[6]. Processes engaged in *CRC* create tightly coupled systems via feedback loops; no boundaries between system components can be isolated. In a much-used example of a *CRC* system, the Watt Governor's arm angle and engine speed are engaged in a co-determining and co-determined relationship[7]. We cannot understand this via a computationalist treatment, since there are no representations to be manipulated via discrete operations in the governing process[8]. Instead, continuous and temporally simultaneous relations of causal influence obtain among the governor's components. No one event can be privileged over the others as the primary determining event: the three way relationship (represented item, representation and user) needed for a representational story is missing[9]

[4]Statements of the computationalist thesis can be found in (Newell and Simon, 1976; Fodor, 1975; Haugeland, 1985; Pylyshyn, 1984, 1987).

[5]Dynamical systems are generally represented by evolutionary (in time) partial differential equations. For an introduction, see (van Gelder and Port, 1995).

[6](Clark, 1998)

[7](Clark, 1998), pp 357-359. Arms on a vertical spindle control steam flow through a throttle valve; the higher the arm angle, the lower the steam flow. The turning of the spindle (the engine speed in other words) generates centrifugal force which causes the arms to fly out. The engine speed causes the arm angle to increase, which reduces steam flow to the engine causing it to slow down and the arm angle to decrease. This then leads to an increase in steam flow leading to an increase in engine speed and so on.

[8]There is scope for disagreement with this as well; Bechtel (1996) presents a contrary view.

[9]One response to the *DH*'s contention that the complexities of certain interactions defy representational treatment is the following (Clark 1998; page 373):

> it is by no means obvious that there are not distinct inner states playing specific representational roles that *may be identified* [my emphasis] by future neuroscientific investigations.

There is an obvious problem with this argument: it is only a promissory note concerning

Computational analysis is also guilty of not solving the *boundary problem*: cognitive systems if dynamical and hence tightly coupled, cannot be understood in terms of an artificial boundary between an inner, privileged, cognitive space and a non-cognitive, purely passive, outer environment. The cognizer and the environment are linked in a complex relationship that resists facile decomposition into temporal and cognitive components. The intrinsic temporality of cognition creates other problems for computational understanding of the mind. Cognition, as a factor in the adaptive success of the organism, takes place in real time: real time involves continuous processes rather than discrete steps. Cognition is *in* time.

DST builds temporal facts into its analysis as each evolving stage of a dynamical system is ordered along a temporal scale. The dynamical analysis for a system coupled with its environment uses evolutionary equations for the coupled system: terms that factor in both systems' current states make up these equations. Such a dynamical model emphasizes the simultaneous, interdependent variation of system components and explains the behavior of a system in terms of alternative ordered sets of changes of the total state over time. This seems to work satisfactorily for the Watt governor and is plausibly viewed as preferable to the alternate computational analysis of the same device[10].

But can the brain be treated as just another dynamical system with no regard to its *inner* structure? The *DC* theorist disagrees[11]:

> the staggering structural complexity and variability of the brain..is the key to understanding the specifically *intelligence-based* route to evolutionary success..that route involves the ability to become apprised of information concerning our surroundings and to use that information as a guide to present and future action.

The primary distinction between the Watt Governor and cognitive systems, it is claimed, is that we do not *need* to give an information based account of the former; an understanding of its operation is achievable by focusing on physical laws relating physical quantities. To maintain this difference, *DC* contends that we need an approach that explains the plausible distinction

the *possibility* that representations *might* be isolated in the future.

[10](van Gelder, 1995), pp 347-350.

[11](Clark, 1998), page 369

between information-based and merely complex, adaptive systems: the *DH* is inadequate as an explanation without the use of representations to make a distinction between characteristics of cognitive systems and just adaptive systems.

One such characteristic is[12]:

> **The Ability to Reason in Absence of Stimulus:** Reasoned behaviors selected in the absence of any local guiding stimulus

That is, cognition is linked with 'intelligent' behavior in the absence of any constant, lawful and reliable signals from the environment. The need for representation arises, then, because cognitive agents are best understood as information-processing systems that *can* reason about environments not immediately available to them. How is this accomplished and what sorts of representations play a role here?

3 The *DC* Thesis

The *DC* thesis has two central, inter-related components:
1. The capacity for partial coupling.
2. The presence of emulators as representations.

So, *DC*'s *prima facie* case begins with the contention that many cognitive capacities are not linked to a tightly coupled relationship with the environment. Cognitive systems can reason about temporally or spatially distant events: the representational approach is claimed to be the most plausible in cases involving reasoning about the distant, the non-existent or the highly abstract.

In these cases, representing and cognizing is achieved by breaking the coupling or the causal loop[13] that exists between the cognitive agent, A and the environment, E. Since the agent half cannot exist independently, we need a modification. We need an executive part, a *controller*, C, for the agent which decouples from E, the environment, and couples instead to E' that is a *stand-in* or representer (or emulator, see below) for the environment.

[12](Haugeland, 1991), (Grush, 1997b)
[13](Grush, 1997b)

E' emulates E and C couples to E' in[14]:

> *exactly the same way as it normally couples to* E [my emphasis].

The appropriate way to analyze the 'representational brain' in this model is to see its activity as the *interdependent* activity of three entities: controllers (executive components of cognitive agents), presentations (these are sensory or perceptual states), and representations (emulators)[15]. A presentation has causal linkage with its target, a representation doesn't. The latter, *strong internal representation*, does not need constant linkage with the outside (it can work in 'stand-alone' mode), but the former, *weak internal representation*, does. A presentation provides information about the target, a representation is a counterfactual presentation (a mock-up representing outer states of affairs). The controller uses both so that our three-way relationship needed for any representational account is complete.

To get around concerns that all that is being proposed is one system standing in for another, Grush (1997a) suggests that the emulator is analyzable into sub-components with representational roles : these represent components of the real system. Its use by the controller is facilitated through selective coupling on the part of the system it services. Neural circuitry for emulators isn't straightforwardly identifiable yet: the evidence so far is indirect[16]. The claim made for emulation is a strong one: identifiable instances of emulator circuitry would be the most foundational and useful way to conceive inner states as representations of the outer environment.[17]

This sort of modeling is best understood as $DST+$ representation through *potential emulation*. Of course, even with E' presumably, there is *strong coupling* present. In each case, the explanation that will be given of the system will be a total state explanation:

> $A \oplus E$ or $A \oplus E'$ where $\oplus$ represents a *tight coupling* operation.

Such emulator modeling comes about because of the benefits it provides; it enables, as stated before, counterfactual reasoning in the absence of environmental stimuli. Actions and goal-achievement strategies can be planned

[14](Grush, 1997b), page 239

[15](Grush, 1997a), page 9

[16](Ghez, 1991)

[17](Clark and Grush, 1999)

'offline' before being actually implemented. Clark and Grush (1999) present the example of motor emulation that facilitates smooth real-time action in the absence of appropriately timed proprioceptive feedback. This enables an evolutionary story to be told about the reason that motor emulation became a neural feature: it enables planning and improves real-world, real-time responsiveness. Emulation is the bridge between the *DH* and computationalism. It is the use of internal representation by emulation that distinguishes cognition from mere adaptive behavior.

The thesis claims to gets its explanatory power from the suggestion that representations via emulation associate purpose to mechanism. In the reaching example, the representation description helps fix the equivalence class of dynamical systems that, as before, fills a specific functional role– that of enabling smooth real-time reaching. Furthermore, seemingly diverse dynamical systems appear united by function with the addition of representation to their analytical treatment.

Most significantly, the emulation condition is a *necessary condition* for a system to be deemed cognitive[18]:

> Cognizers, then, use models (internal and external) in place of directly operating upon the world. Non-cognizers by contrast, remain trapped in a (potentially very complex and context-variable) web of closed -loop interactions with the very aspects of reality upon which their survival depends.

This condition, unfortunately, undermines the plausibility of the *DC* thesis.

4 Problems with *DC*

The *DC* claim carries empirical plausibility since it makes a case for some evolutionary advantage for representation via emulations. I do not doubt that emulation does take place in *some* cognitive systems[19]. However, while the premises of the *DC* argument might be true (that a distinction needs to be

[18] (Clark and Grush, 1999)

[19] The empirical evidence for the existence of emulator systems in cognitive systems is presented in (Wolpert, 1995), (Kawato, 1990).

made between cognitive and adaptive systems, for example), its conclusion that the capacity for emulation is a necessary condition for cognition, is problematic, if not just false. There are several problems with the model and I'll consider each one in turn.

Firstly, as a description of our cognitive activities (or the activities of other systems whose classification as cognitive meets with some consensus) in terms of either frequency or importance, the ability to reason without local stimulus present is neither a complete nor an accurate description. A large percentage of our actions consist of bodily actions guided by mind (in so smooth and automatic fashion so as to make them almost thoughtless) that in fact *do not* consist of reasoning about the distant, non-existent, or the highly abstract. The importance of this is simple: we can easily imagine cognitive agents that do not reason about the distant, non-existent or highly abstract; animals being a good example. When they are not doing so, they are not un-cognitive, they simply happen to be cognizing in a different fashion. Any theory that suggests that one kind of cognition is more important than the other is going to have to tell a fairly detailed story about the evolutionary benefits it entails. Perhaps one can be told; but whatever its form, it should not have as a consequence that systems that do not have this ability cannot be deemed cognitive. This ability can be a distinguishing characteristic only when the systems being excluded also display a gross lack of behavioral flexibility (and perhaps other features as well).

Secondly, to stipulate the existence of a controller in the tripartite representational model of the brain is to beg the question, since it is the existence of such a controller that is being disputed by the *DH*. To postulate such an entity is to give a homuncular treatment of the brain; precisely the sort of thing the *DH* disputes. Even if the controller were to be admitted, how can it be identified? None of the discussion makes any nod in the direction of how it would be possible to empirically identify the *part* of the agent deemed the controller. If an analysis is to be couched in dynamical terms as the *DC* theorist admits, how is one to isolate the entity termed the controller?

Thirdly: This partial or selective coupling leads to the ability to cognize[20]:

> even when not *actively engaged* with a target system or interactive environment

[20](Grush, 1997a), page 12

Here, a natural language confusion over the meaning of the word 'actively', obscures the point of the *DC* theorist. It's not clear what *not being actively engaged* with the environment amounts to. There is *always* some environment that you are *engaged* with, otherwise I wouldn't be conscious or aware at all of my surroundings (an agent 'engaged' only with an inner environment implies either unconsciousness or some form of meditation.) I might reason about my European vacation while working in the office, but at any point, I am in contact with my environment, aware of being in such a state and potentially capable of taking action on receipt of input from the environment. Does not being actively engaged with the environment mean that I have disengaged to the extent that my environment has become irrelevant? Does it mean that none of my actions can be traceable to inputs from the environment?

The distinction between being actively engaged and not being so engaged with the environment is unclear: the agent is always engaged in constant interaction with *some* part of the environment. Reasoning about the absent is just one of the *multiple* tasks taken on by the agent in the course of its cognizing. One example of not being 'actively' engaged with an environment might be suspension in a sensory-deprivation tank. In any other case, it isn't immediately apparent how the agent can be said to be disengaged from the environment. Perhaps the agent is engaged partially with one environment and partially with another? If that is what is intended, then the implication that the agent 'detaches' itself from some privileged, isolated environment, E, to couple to another, E', is confusing. There may be more than one coupling in progress at any given time, but the difficulties of isolating coupled pairs in a cognitive system should convince us that we would be better off viewing the evolution and behavior of the system as a whole and providing a total state explanation instead.

Fourthly: has selective coupling solved the problem created by tight coupling? What it has done is is replace it with a *varieties of coupling* thesis. At no point, is there *no coupling*; one can assume that the coupling represented by $A \oplus E'$ is a tight one; and in fact, it is:[21]

> Provided that the emulator is accurate, and that *its* output is not delayed, all will proceed without a hitch.

[21] (Grush, 1997a), page 13

Timing issues still remain critical: A must decouple, couple and then couple back again in order to take action. The sequencing done in real-time of this constant coupling and decoupling is critical to the success of the organism. The *DC* theorist, then, has postponed the problem one step further. Essentially, we have replaced the tight coupling of $A \oplus E$ with $C \oplus E'$. How, then, has the *DC* solved the problem of what Grush terms 'obligatory action'? If the controller is still tightly coupled to the emulator during its interactions then there isn't a stage when a controller isn't in a tightly coupled loop.

The *DC* is guilty, too, of simply conflating 'intentional' with 'representational'. As an example of the *DC*, Grush presents Murphy, a robotic system (initially described in (Mel, 1990)) that uses a visual grid as an emulation system for 'off-line' trials of potential grasping moves before actually implementing a solution. Grush claims that any description of its behavior will be couched in representational terms. The explanation given, though is simply *intentional*[22]:

> Why did Murphy not move its arm to location X? Because it..(*believed*) [my emphasis] that if it did, its arm would hit an obstacle. What is Murphy thinking about now? It is *contemplating* [my emphasis] the use of command sequence R...some sort of representational description is almost forced by Murphy's use of the emulator to stand for the real system.

This renders 'intentional' *synonymous* with 'representational'. But there is no reason to think that such a synonymy obtains. All this shows, instead, is that one can, to co-opt a bit of standard phraseology, adopt an *intentional stance* towards Murphy. The explanation above, instead, sounds circular; a representational system is intentional, indulges in reason-driven action, and an intentional system is one that uses representations to aid in its task of information processing. The claim, then, that an intentional system is in fact, using *representations* still needs to be argued for[23].

Allied to the varieties of coupling problem is the fact that all the emulator theory has accomplished is move the environment *inward*. It still isn't clear where the representations are. In an attempt to get rid of 'ontologically

[22] (Grush, 1997a)

[23] See (Dennett, 1987) for descriptions of difficulties in making conceptual associations between representations and intentional systems.

magic barriers'[24], it has suggested a barrier between the controller and the emulator. Where does one begin and the other end? Why is this any different from suggesting that the skull is an 'ontologically magic barrier'? Given that the environment can now plausibly be said to consist of E and E', in that one can still make a distinction between controller + agent and everything else, how has the boundary problem been addressed?

The most serious problem however, is in the elevation of the capacity for emulation to a necessary condition for cognition.

4.1 The Value (Necessity?) Of Representation by Emulation

Controllers without emulators do not represent the events they control, they always need physical feedback from extra-neural states of affairs. This lack of a dependence, then, for *DC*, is the distinguishing characteristic of cognitive systems. *No cognition without representation via emulation.* But, does this then mean that if there is no capacity for representation via emulation, there is *no* cognition? Apparently so. That a complex, adaptive system needs to do this in order to acquire *added* evolutionary success is a claim that could be argued for. Its elevation to a necessary condition, however, is more problematic. But elevated it is, as follows[25]:

> **The Emulation Condition:** A system is a representing, cognitive system only if it has the capacity to internally emulate some external system with which it can also interact overtly

One problem in thinking of terms of capacities being necessary conditions for cognition is that it leaves open the possibility that a system, plausibly viewed as cognitive on other counts, could have a 'capacity' to decouple and emulate but never actually do so. In that case, the questions that can be asked of the *DC* theorist are numerous: What plays the part of representations? The unused emulator? Is such a capability *necessarily* invoked? If a system never actually uses the emulation 'capacity', then in what sense does it possess it? Could we ever point to a piece of neural circuitry and call it an emulator without that piece having been used as such? Consider

[24] (Grush, 1997b), page 239
[25] (Grush, 1997a)

two adaptive systems, identical in behavior (right down to histories), but possessing different physical organizations. One system possesses what looks like an 'emulator' and one doesn't. Both do not actually carry out emulations. What, at this state warrants our judging one to be cognitive, and the other to be 'merely' adaptive? It's plausible that emulation systems make their appearance when it's found that real-world responses improve as a result, but then, those systems that started to use emulation as a success strategy didn't *become* cognitive with the addition of the emulation capacity.

Furthermore, it's plausible that the capacity to carry out emulation could be met by a diversity of different physical organizations (multiple realizability for emulation). This creates the problem, then, for the *DC* thesis that the restriction is not of much use in actually isolating a cognitive system from other adaptive systems, since the two could be reasonably similar in physical structure and thereby display the same 'capacities' without us granting the same status to both systems. Furthermore, what about emulator circuitry makes it so special that an inspection of it could reveal the capacity to emulate? If an evolutionary story about the appearance of emulation on the cognitive scene is to be told, presumably this will be along the lines of a pre-existent physical feature being pressed into service for the purposes of emulation. One would distinguish it from other neural circuitry only at this point, i.e. at the point of usage.

Since the presence of a capacity can only be revealed by that capacity being exercised, why not couch the condition in terms of the capacity being used? That is: 'an adaptive system is a cognitive system only if it carries out emulations of external states of affairs'? For the very good reason that it would be an improbably strong condition. It's possible that some systems do not reason about the distant, the absent (and do not use or need the capacity for emulation) and still manage to make perfectly good if not spectacular cognitive agents: the absence of this emulation makes no difference to their evolutionary success. The *DC* model cannot rule out the possibility of such agents. The necessity of this condition is suggested, simply because, I suspect, it is felt to be a plausible description of how some kinds of cognitive agents do function in parts of their lives. But this, then, fails to be a necessary condition. The *DC* explanation is more plausibly viewed as a description of a subset of possible cognitive systems.

The virtue of this approach is held to be that it is less liberal than dynamical systems approaches in that it precludes a class of systems. But it might have precluded too much. The notion of a capacity, as noted, is only coherent, when it is actually used. Call the possession of the capacity for emulation a *conceptual constraint* on the class of cognizers. Is this constraint an intuition about cognizers? If so, how does it get its grounding? Is this a stipulation or an empirical claim about cognizers? If an empirical claim, its lack of empirical support is disturbing. Placing a constraint like this on the class of cognizers, especially as a necessary condition is a move that needs some explaining.

4.2 Defending Dynamical Computationalism

What sort of defense against the criticisms above can be mounted by the *DC* theorist? Consider the following[26]:

Defense 1: 'a hammer is a hammer even if it's never used. One would still want to make a distinction between a hammer and a bag of chips in terms of their ability to drive nails, even if never used. More importantly, a creature that has the capacity for emulation, but never used it presents no more than a taxonomic problem.'

Response: In the case of a hammer, we know what it is: recognizing a hammer's function takes place in the context of acquaintance with other members of the hammer class. We know it's a hammer, a device used for driving in nails. It has the capacity for driving in nails (even if never used) by virtue of its membership in the hammer class. It isn't clear, though, that a hammer's function is that of knocking in nails if we have *no prior* experience with members of the hammer class (if hammers were found objects, for example). In the case of a cognitive system, judgement of cognitive status takes place at the species level; if the species displays cognitive behavior and some of its members show the capacity for emulation by actually employing it, then we will judge other members to have the capacity, even if they don't use it. If we are at the stage of initial judgement, and no members of the species use the 'capacity', then the 'capacity', without actual usage, cannot be meaningfully detected or pinpointed. Potential usages will not pick out some unique usage, but instead some set of possible uses. This renders the emulation thesis too vague to be of any real help in judging cognitive status. As for being 'just a taxonomic problem', well, the discussion is about classification of different

[26]The first three of these defenses are due to Rick Grush (personal communication).

kinds of systems, so it can't be anything but a taxonomic problem–just one indicative of a very real difficulty with the basis for the taxonomy.

Defense 2: 'The more interesting issue is : is a creature that never cognizes anything a cognizer if it *could* cognize things? As long as we are clear on what is at issue, not much hinges on the answer to this question. There is no problem in this case in saying that such a creature is not representing or cognizing anything.'

Response: This isn't a problem because the case being addressed occurs when a creature displaying behavior that we could plausibly term cognitive, is not classified as such because of its failure to possess an unused 'capacity' for emulation. If a creature never cognizes, but has the 'capacity' to cognize (once again, I'm not sure how this capacity could be determined) then I agree this creature is not a cognizer, but that is to concede my point, not dispute it.

Defense 3: 'What is at issue is tight coupling with the environment and other kinds of coupling don't matter'.

Response: It isn't clear however, why the emulator system shouldn't be considered part of the environment. So long as we make a distinction between the agent + controller combine and everything else, it isn't clear why the environment doesn't just extend inward.

Defense 4:[27] Anti-representationalist rhetoric focuses on the complexities of closed loop agent-environment interactions that show a particularly intractable form of *CRC* not amenable to a representational treatment; the *DC* thesis is focused on cases where the agent decouples.

Response: If the result of the decoupling is an engagement is another tight coupling, then, how are we any closer to a representational treatment of the kind desired? Furthermore, the decoupling strategy itself comes up as a response to the *DH*'s contention that closed loop interactions are exhaustive descriptions of agent environment interactions.

Defense 5: A representational story is the best explanation[28]:

> The best criteria indeed for the legitimacy of a representational explanation are whether or not such explanations make sense of the behavior of the system.

[27](Clark and Grush (1999))
[28](Clark and Grush, 1999)

Response: To this, I would add the following caveat: 'and provided there are no better explanations forthcoming'. Is the representational story told by the emulation thesis so good that it must be told for every conceivable cognitive system? That a system for which it cannot be told fails to be a cognitive system? If a creature were to display flexibility, the ability to evaluate different responses 'on-line' and a history of evolutionary success, but no capacity to emulate, then it seems implausible that this absence is the only reason to disqualify it from cognitive status. Some additional disqualification is required.

5 Conclusions

Even if we are to count emulations as representations, temporal dynamics must be attended to in order to do justice to real-time constraints. So, we need representational pictures and dynamical systems analysis: dynamical modeling can, should, and must play a part. The *DC* theorist concedes that it's possible that a purely dynamical picture will land up being the correct one, since it seems that computationalism has made the following adjustments in its putative reconciliation with the dynamical hypothesis:
1. Discrete, individuable states that are the subject of computation and that form the constituents of mental transitions have been abandoned in favor of a dynamical approach.
2. The *CRC* nature of cognition has been acknowledged.

DC holds on to the idea that representation is necessary, and indeed, representation by emulation is a necessary condition for a complex, adaptive system to be considered cognitive. Its primary problems are the use of the unclear notion of not being actively engaged with the environment and the necessity of a condition that precludes large classes of agents that we plausibly deem cognitive. Agents that do not decouple ever from their environment and yet, are still capable of highly purposive action cannot be ruled out by such a treatment.

Emulators, if ever conclusively identified and functionally tagged, could be termed representations, much like other representations proposed by the *DH*; their absence needn't be a condition for ruling out cognitive status to a possibly cognitive system.

Acknowledgements: I'd like to thank Andy Clark, Tim Van Gelder and

Rick Grush for helpful comments on an earlier version of this paper.

References

Bechtel, William. (1996). Yet Another Revolution? Defusing the Dynamical System Theorists' Attack on Mental Representations. *Presidential Address to Society of Philosophy and Psychology.*

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

Clark, Andy and Grush, Rick. (1999). Towards a Cognitive Robotics. *Adaptive Behavior*, Vol 7:1.

Dennett, Daniel. (1987). *The Intentional Stance*, MIT Press, Bradford Books.

Fodor, Jerry. (1975). *The Language of Thought.* New York, Crowell.

Ghez, C. (1990). Voluntary Movement. In Kandel, E., Schwartz, J., and Jessell, T. eds. *Principles of Neural Science*, Elsevier, Holland.

Grush, Rick. (1997a). An Architecture Of Representation. *Philosophical Psychology*, 10(1) 5- 23.

Grush, Rick. (1997b). Yet Another Design for the Brain? Review of *Mind as Motion. Philosophical Psychology*, 10(2) 233-242.

Haugeland, J. (1985). *Artificial Intelligence: The Very Idea.* Cambridge, MA: MIT Press.

Haugeland, J. (1991). Semantic Engines: An Introduction to Mind Design. In J. Haugeland (ed.) *Mind Design: Philosophy, Psychology, Artificial Intelligence*, pp 1-34, MIT Press, Cambridge, MA.

Kawato, M. (1990). Computational Schemes and Neural Network Models for Formation and Control of Multi-joint Arm Trajectory. In W.T Miller III, R. Sutton and P. Werbos (eds.), *Neural Networks for Control.* Cambridge, MA: MIT Press.

Mel, Bartlett. (1990). Vision-based Robot Motion Planning. In W. Thomas Miller III, Richard S. Sutton and Paul Werbos (eds.) *Neural Networks for Control*, MIT/Bradford Books.

Newell, A., and Simon, H. (1976). Computer Science as Empirical Enquiry: Symbols and Search. *Communications of the Association for Computing Machinery*, 19, pp 113-126.

Pylyshyn, Z. (1984). *Computation and Cognition: Toward a Foundation for Cognitive Science.* Cambridge, MA: Bradford/MIT Press.

Pylyshyn, Z. (1987). (Ed.) *The Robot's Dilemma: The Frame Problem in Artificial Intelligence*, Norwood: Ablex. 1987.

Thelen, Esther and Smith, Linda B. (1994). *A Dynamic Systems Approach to the Development of Cognition and Action.* Cambridge, MA: MIT Press.

Van Gelder, T. and Port, Robert. (1995). It's About Time: An Overview of the Dynamical Approach to Cognition. In van Gelder, Tim, and Port Robert, (eds.) *Mind as Motion: Explorations in the Dynamics of Cognition*, MIT/Bradford.

Van Gelder, Tim. (1995). What Might Cognition Be, If Not Computation? *Journal of Philosophy*, XCII, 7, July 1995.

Wolpert, Daniel, Ghahramani, Zoubin and Jordan, Micheal. (1995). An Internal Model for Sensorimotor Integration. *Science* 269: 1880-1882.