

University of Wollongong

Research Online

Faculty of Law, Humanities and the Arts -
Papers

Faculty of Arts, Social Sciences & Humanities

1-1-2013

Cognitive assembly: towards a diachronic conception of composition

Michael D. Kirchhoff

Macquarie University, kirchhof@uow.edu.au

Follow this and additional works at: <https://ro.uow.edu.au/lhapapers>



Part of the [Arts and Humanities Commons](#), and the [Law Commons](#)

Recommended Citation

Kirchhoff, Michael D., "Cognitive assembly: towards a diachronic conception of composition" (2013).

Faculty of Law, Humanities and the Arts - Papers. 1241.

<https://ro.uow.edu.au/lhapapers/1241>

Research Online is the open access institutional repository for the University of Wollongong. For further information contact the UOW Library: research-pubs@uow.edu.au

Cognitive assembly: towards a diachronic conception of composition

Abstract

In this paper, I focus on a recent debate in extended cognition known as “cognitive assembly” and how cognitive assembly shares a certain kinship with the special composition question advanced in analytical metaphysics. Both the debate about cognitive assembly and the special composition question ask about the circumstances under which entities (broadly construed) compose or assemble another entity. The paper argues for two points. The first point is that insofar as the metaphysics of composition presupposes that composition is a synchronic relation of dependence, then that presupposition is inconsistent with the temporal dynamics inherent in the process of cognitive assembly. The second point is that by developing a diachronic or temporally dynamic ontology for understanding the phenomenon of cognitive assembly, this lends support for a third wave of extended cognition.

Keywords

conception, assembly, cognitive, composition, diachronic, towards

Disciplines

Arts and Humanities | Law

Publication Details

Kirchhoff, M. D. (2015). Cognitive assembly: towards a diachronic conception of composition. *Phenomenology and the Cognitive Sciences*, 14 (1), 33-53.

Cognitive Assembly: Towards a Diachronic Conception of Composition

Michael D. Kirchhoff

Please do not cite, forthcoming in Phenomenology and the Cognitive Sciences

1. The extended cognition thesis

The target of this paper is the extended cognition thesis (EC). According to EC, in certain circumstances, many of our cognitive processes and cognitive systems are distributed not only across various regions of the brain but also across non-neural bodily and environmental parts (Clark 1997, 2008; Clark & Chalmers 1998; Menary 2007; Sutton 2010; Wheeler 2010).

When addressing questions and issues pertinent to EC, there are at least two explanatory targets one must be careful to distinguish between. The first is the assembly process involved in orchestrating an extended cognitive process or an extended cognitive system. I refer to this as *the process of cognitive assembly* (the process of CA). Notice that when addressing the process of CA, we may ask over which timescales, and by which principles, different entities (e.g., certain neural operations, saccadic eye movements, cultural practices, tools) jointly combine to put together or compose a temporally distributed cognitive process and/or system. The other target concerns the newly assembled device once it has been assembled or composed. I call this *the product of cognitive assembly* (the product of CA). In relation to the product of CA, we may ask questions about the difference in properties between the various parts, how information flows and its propagated between the different parts, and how it is that an extended cognitive system instantiates such “processes in ways that ideally solve some problem.” (Clark 2008, p. 122)

In this paper, I focus on a recent debate in EC between Clark (2008, 2011) and Hutchins (2011a) concerning the *process of CA*. My independent motivation for this are twofold: first, to sort out some of the positions one might take in this discussion between Clark and Hutchins; and second to defend that at least one of those choices marks a step towards a third-wave version of EC.

2. Towards a third wave of EC: Setting up the arguments

When addressing the process of CA, what I call the usual account is as follows: Only processes operating in the here-and-now are responsible for the process of CA and such processes are primarily bodily and neural processes. Prominent advocates of this view include Clark (2008, 2011) and Clark & Chalmers (1998). I should add that the usual account of EC is what has recently been referred to as *first-wave EC*. Defenders of first wave EC ground their arguments for EC on the ideas that (i) external artifacts are incorporated into the cognitive system of an individual in virtue of the right kind of causal coupling, and (ii) functional similarity between the causal roles of internal and external occupiers. Alternatively to first-wave EC, defenders of *second-wave EC* go beyond parity and focus instead on complementarity between internal and external states and properties (Sutton 2010) and their consequent integration into a cognitive whole (Menary 2007).

In this paper, I explore one possible route by which to gesture at but also argue for a third wave version of EC. Specifically, I follow work by second-wave EC theorists, whom are driving particular visions for a third wave of EC. In particular, I propose one approach to the process of CA that exemplifies Sutton's recent gesture towards a third wave of EC: a version "which dissolves individuals into peculiar loci of coordination and coalescence among multiple structured media [and practices]." (2010, p. 213) This suggestion, I believe, echoes what Menary has recently called "enculturated cognition" (EnC). EnC is the "idea that our cognitive abilities are transformed by a cognitive species of cultural practices [...]. What we are able to do is augmented and transformed by the acquisition of cognitive practices." (2012, p. 148) Both of these versions for a third-wave of EC theorizing emphasize the deconstruction of the individual organism as the locus of the process of CA and allow for cultural practices as playing a central role in close coordination with neural and bodily processes.

In this paper, I aim to unpack this articulation of a third wave in the extended cognition thesis in a way that has not been done before. That is, I will discuss the process of CA in conjunction with work on the relation of composition in metaphysics. I should add that even though the debate about the process of CA is not strictly speaking about the metaphysics of composition, those involved in the discussions over CA ask structurally similar questions to those involved in debates

about composition in metaphysics. In debates over composition, what is known as the *special composition question* (SCQ) (van Inwagen 1990) concerns the circumstances under which entities assemble or compose another entity. Indeed, both the process of CA and SCQ ask questions concerning the conditions under which entities combine to compose or assemble another entity (or whole). In so doing, both the SCQ and the process of CA take as their target the Xs – the constituents – that compose or assemble Y, and analyze the conditions under which the Xs come together to compose or assemble Y.

The first argument that I will develop turns on the fact that composition is understood as a synchronic relation of dependence, and that such a synchronic notion of composition (the SCQ) is inconsistent with the temporal dynamics inherent in the process of CA. To make this claim, I aim to establish that the restriction of the verb “compose” in the expression “the Xs compose Y” to the *present tense* is metaphysically problematic when considering the nature of time continuous processes such as those involved in the process of CA. This picture, familiar as it is, of X (or the Xs) composing Y at an instant *t*, finds no corresponding image in contemporary debates about the process of CA. In fact, when Clark states that his own targets in *Supersizing the Mind* (2008) are processes operating in the *here-and-now*, nowhere does Clark’s temporal quantification “here-and-now” adhere to the assumption that the verb “compose” must be understood to imply “compose at an instant *t*”. Consequently, serious inquiry into the kinship between the process of CA and the SCQ must begin by scrutinizing the actual meaning of the term “now” as it is used to express the claim “the Xs compose Y *now* or in the *here-and-now*”. What I shall argue is that the process of CA *must* be stated without implicating a notion of composing or assembling that is synchronic, where “synchronic” means that the Xs are composing Y at this very instant *t*. This is important, since I will show that only by problematizing the notion of synchronic composition is it possible to provide a properly motivated answer to the process of CA. That is where the metaphysical action lies in this paper¹.

¹ What I do not claim is that arguments for the process of CA assume a synchronic conception of composition or assembly. It is the usual accounts of composition in analytical metaphysics that presuppose that composition holds in a synchronic manner. What I do claim is that scrutinizing the meaning of “now” both in the debate

The second argument considers the debate between Clark and Hutchins on the process of CA, with the aim of establishing that as soon as we leave room for the non-trivial role of cultural practices in the process of CA – even when the processes unfold in the here-and-now – this requires that we must look beyond the system made up of the individual agent and artifact. That is, we must include into our explanation of the process of CA features such as cognitive norms (Menary 2007) and patterned (cultural) practices (Hutchins 2011; Menary 2007; Roepstorff et al. 2010). By the end of the paper what I hope to have shown is that developing a diachronic or temporally dynamic ontology for the process of CA lends support for a third wave of EC.

3. Overview

The structure of the paper is as follows. In section 4, I consider difficulties with giving a satisfactory answer to the qualification that the verb “compose” is to be understood in the present tense – i.e., as *right now* some Xs are composing or assembling some Y. In section 5, I make use of the outcome of my discussion in section 4 to discuss one of Clark’s examples of the process of CA, namely Gray & Fu’s (2004) studies on the soft-assembly of interactive microstrategies employed by the brain to solve a given problem. In section 6, I consider the debate between Hutchins and Clark on the process of CA, where I draw up the battle lines and discuss the possibility that Clark might be wrong in privileging bodily and neural processes when explaining the process of CA². In the final section, I tease out several implications of the discussion in this paper.

over the process of CA and in analytical metaphysics may shed light on how (a) to understand the meaning of “now” in debates concerning composition, and (b) to further develop the project of establishing a framework for a third-wave of EC.

² Some readers may find it a controversial claim to state that Clark privileges bodily and neural processes when explaining the process of CA. However, in his most recent book, *Supersizing the Mind* (2008), Clark endorses the following position when discussing the process of CA: “Human cognitive processing (sometimes) literally extends into the environment surrounding the organism. But the organism (and within the organism, the brain/CNS) remains the core and currently most active element. Cognition is organism centered even when it is not organism bound.” (Clark 2008, p.

4. Discussing the terms “now”, “right now” and “here-and-now”³

Under what circumstances do a collection of entities compose some further entity? As we saw in section 2, this is van Inwagen’s SCQ (1990). Some answer never (Rosen & Dorr 2002), some say sometimes, yet only sometimes (Markosian 1998; Merricks 2001; van Inwagen 1990), whereas some say always (Lewis 1986; Sider 2001). In this section, I do not consider any of these options for when (or if) composition holds. Instead, I start by considering difficulties with providing a satisfactory answer to the assumption that the verb “compose” in the expression “the Xs jointly combine to compose Y” is to be understood as meaning “*right now*, the Xs jointly combine to compose Y”.

Consider, for example, what van Inwagen writes about *tense* and *composition* in *Material Beings*, as he says:

“The verb ‘compose’ in the predicate ‘the *x*s compose *y*’ is to be understood as being in the present tense, and the same point applies to ‘are’ in ‘are parts of’. Thus, ‘are parts of’ and ‘compose’ should be read ‘are *now* parts of’ and ‘*now* compose’. Strictly speaking [...], our *definiendum* should have been ‘the *x*s compose *y* at *t*’, and our “primitive” mereological predicate should have been ‘*x* is a part of *y* at *t*.’ (1990, p. 29; italics in original)

139) Or, as Clark puts the point a few pages earlier: “It is indeed primarily (though not solely) the biological organism that, courtesy especially of its potent neural apparatus, spins and maintains (or more minimally selects and exploits) the webs of additional structure that then form parts of the machinery that accomplishes its own cognizing.” (2008, p. 123) These two quotes are indicative of Clark’s *Hypothesis of Organism Centered Cognition*, and is what I mean by the claim that Clark privileges bodily and neural processes when addressing the *process of CA*. But, of course, this does not imply that the Hypothesis of Organism Centered Cognition is inconsistent with EC. Indeed, Clark’s position on the product of CA is precisely that certain systems are distributed across neural, bodily, and environmental resources.

³ I use the terms “now”, “right now” and “here-and-now” interchangeably throughout this paper.

On this view, for the Xs to compose Y *now* is to claim that the Xs compose Y at *t*. By my lights, this is a rather elusive way to characterize the idea of “now” in relation to composition, because it leaves it an open question whether van Inwagen means either (a) that composition is a relation between the Xs and Y at time *t*, where the extension of *t* may include a short interval of time, or (b) that composition is a relation between the Xs and Y at a momentary instant, where the extension of *t* does not include an interval of time but denotes a durationless point in time. However, consider what Bennett says about composition in her survey of metaphysical dependence relations: “*Composition* is a synchronic or atemporal many-one relation [...]” (2011, p. 81). Furthermore, the usual account of composition depicts composition as a *vertical relation of dependence* between Xs and Y, where ‘vertical determination’ must be understood to exclude ‘horizontal determination’ (Bennett, forthcoming; Kim 2005). Put differently, if there is a relation of vertical determination between some Xs and a Y, then that relation of determination does not unfold across an interval of time – where horizontal determination represents diachronic (temporal) causal relations. The usual account of composition – in metaphysics as well as the metaphysics of science – turns on the presupposition that composition is a relation that holds at a durationless instant.

Prima facie, at least, it seems to me that ordinary folk are quite familiar with the idea of “now” or “right now”. For instance, when we talk about something *presently* taking place such as executing a tennis serve, engaging in a conversation with a friend, stirring the pasta sauce, and so on, we (implicitly) appeal to the temporal fact that something – which we may or may not engage in – is happening or taking place *right now*. If the folk are right, which they may or may not be, it would make the set of events to which the term “now” applies include events the temporal duration of which takes place over milliseconds, seconds, and minutes. It may only take me a few seconds to stir the pasta sauce, but it may take me several minutes to execute the proper chess move *vis-à-vis* the situation on the chessboard⁴.

In metaphysics things are different from that of the folk. To facilitate our discussion, consider that Markosian (2004) mentions that it is possible to distinguish between two different senses of the notion “X exists now”. The first sense is what

⁴ I should point out the duration of “now” is also discussed in the philosophy of time consciousness under the heading of the *specious present*.

Markosian calls for the *temporal location* sense, where the expression “X exists now” is meant to be synonymous with “X is present”. This is the received view of *presentism* in the philosophy of time (Markosian 2004). The second sense of “X exists now”, Markosian calls for the *ontological* sense. On this ontological sense, the expression “X exists now” is understood as shorthand for the claim that X “is now in the domain of our most unrestricted quantifiers, whether it is present, like you and me, or non-present, like Socrates.” (Markosian 2004, p. 48) This is the view commonly expressed as *non-presentism* or *eternalism* (Sider 2001). Presentism and non-presentism make competing claims about temporal ontology (Sider 2006). Presentism is the view that, necessarily, it is always true that only present entities exist, whereas non-presentism is often formulated in an atemporal language that is hostile to presentism (Sider 2001). That is, on the non-presentist view, past and future entities, such as dinosaurs and me 10 years from today, all exist. However, for the presentist, but not the non-presentist, there is something ontologically special about the *now*, in the sense that *only* entities that are currently present, exist.

For my purposes here, it matters little whether presentism or non-presentism is ultimately true, in that, all I wish to highlight is that both of these (hotly debated) doctrines in metaphysics are equally elusive when it comes to pinning down the precise meaning of the notion “now” as the above given definition by van Inwagen (1990). First, the presentist states that only present entities exist. But what might this claim amount to in the context of cognitive science? Consider, for instance, the Hodgkin & Huxley (1952) model of the action potential discussed in Craver (2007).

Action potentials consist of both rapid and fleeting changes in what is known as the electrical potential difference in a neuron’s membrane. This electrical potential difference (measured as the voltage difference across the membrane) is known as the membrane potential. The membrane potential consists of a separation of charged ions on either side of the membrane (*cf.* Craver 2007, p. 50). As Craver specifies: “In the neuron’s resting state, positive ions line up against the extracellular surface. In typical cells, this arrangement establishes a polarized resting potential (V_{rest}) of -60 mV to -70 mV [...]. In an action potential, the membrane becomes fleetingly permeable to sodium (Na^+) and potassium (K^+). This allows the ions to diffuse rapidly across the cell membrane.” (2007, p. 50) This rapid diffusion changes the mV, in that, the action potential consists of (i) a quick increase in mV to a maximum of $+35$ mV, which is followed by (ii) a rapid decrease in mV to certain values below the so-called V_{rest} ,

followed by (iii) a prolonged or extended after-potential period during which the neuron is less excitable (*cf.* Craver 2007, p. 50)

Unconventional as this example might be in discussions over presentism, the question I wish to consider goes to the very heart of the ambiguity of how to understand “present” or “now”. That is, is the action potential *present* in the presentist sense? Let us start with the (arguably) uncontroversial assumption that insofar as an action potential is present now (whatever we take “now” and “present” to imply), the action potential requires for it to *be* present that it *unfolds over a region of space-time*. That this is uncontroversial follows from the brute fact that an action potential irreducibly consists of (i), (ii), and (iii) – that is, of (i) a quick increase in mV to a maximum of +35 mV, (ii) a rapid decrease in mV to certain values below the so-called V_{rest} , and (iii) a prolonged or extended after-potential period during which the neuron is less excitable (*cf.* Craver 2007, p. 50) – and the manifestation of any of these three stages or parts *take time* (measured in milliseconds). The presentist is committed to the ontological claim that only present entities exist. However, consider some event – the rapid rise in mV to a maximum value of approximately +35 mV – that is happening right now. Too late! That event is over, in the sense that the mV is already rapidly declining to a value below V_{rest} . If we take seriously that it is only present entities that exist, it follows that the first stage of the action potential is now entirely in the *past*. However, according to the presentist, everything that is either in the past or in the future (or both) does not strictly speaking exist; only entities that are present exist *now*.

I do not intend this to be a refutation of presentism; however, if this is indeed one possible outcome of presentism, it gives rise to a counter-intuitive situation: that the first stage that the action potential consists of is *not* – or, on longer – part of the action potential because of the fact that the mV is *presently* on the decline. Furthermore, what is *now* part of the action potential – for example, stage two – will, in a very short period of time, cease to be a part of the action potential because it will be entirely in the past. For the cognitive scientist as well as philosophers of cognitive science this result, I suspect, will be unbelievable. For example, in mechanistic philosophy of cognitive science (Bechtel 2008; Craver 2007; Machamer et al. 2000), lower-level components and their activity give rise to higher-level functioning in virtue of the components being organized in a certain temporal, spatial as well as causal organization. But, the components themselves will not cease to be part of some

higher-level phenomenon or mechanism, because they operate over different temporal frequencies.

Non-presentism will do no better for my purposes, especially because non-presentism is stated *tenselessly* or *timelessly*, thus completely ignoring one of the central principles of research across cognitive neuroscience (Engel et al. 2001), cognitive psychology (Ballard et al. 1997; Spivey 2007) as well as dynamical and embodied approaches to cognition (Clark 1997; Gibson 1979; van Gelder 1998; Varela et al. 1991; Wheeler 2005; and others), namely that cognition *happens in time* (not in some atemporal or timeless vacuum), and that time *constrains* and *limits* the production of cognitive activity. Wheeler captures this emphasis on temporality nicely, as he states:

“In the psychological arena, such phenomena [i.e., temporally rich phenomena] include (i) the rates of change within, the actual temporal duration of, and any rhythmic properties exhibited by, individual cognitive processes, and (ii) the ways in which those rates change, temporal durations, and rhythms are synchronized both with the corresponding temporal phenomena exhibited by other cognitive processes, and the temporal processes taking place in the cognizer’s body and her environment.” (2005, p. 106)

The central problem with both the standard synchronic notion of composition, on the one hand, and the doctrines of presentism and non-presentism, on the other, is that none make concessions to the fact that time is *continuous* (Spivey 2007): one that impedes the treatment of time in terms of arbitrary, discrete step time (t_1 , t_2 , etc.).

The standard composition theorist, however, is a *synchronic* composition theorist in that he/she accepts that the primitive “X exists at time (or temporal interval) t ” implies ontological synchronicity with respect to t . That is, the synchronic composition theorist claims that if X exists at t and if X is part of Y, then X is part of Y at t – this will be so no matter how continuous or discontinuous, transient or durable the interval and sub-intervals may be.

This practice of casting the temporal conditions under which X (or the Xs) composes Y into some lockstep or stepwise progression (t_1 , t_2 , etc.), involving a sequence of discrete states – such that X_1 composes Y_1 at t_1 , X_2 composes Y_2 at t_2 , and so on, until t_n – highlights an important difference between temporally complex

forms of composition (as in the process of CA) and the kind of composition most metaphysicians have in mind. To further highlight this difference, consider the following example of returning a tennis serve, in van Gelder & Port:

“The ball is approaching; you are perceiving its approach, are aware of the other player’s movements, are considering the best strategy for the return, and are shifting into position to play the stroke. *All this is happening at the same time.* As you move into place, your perspective on the approaching ball is changing and hence so is activity on your retina and in your visual system ... the path of the approaching ball affects which strategy would be best and how you move. *Everything is simultaneously affecting everything else.*” (1995, p. 23; italics in original)

In this example, claiming that a system S instantiates Y (returning a tennis serve), and Y is composed of some particular Xs, at a particular point in time *t*, really boils down to saying that *during that period of time* Y was composed of the Xs (or, Y was composed by the Xs). Keep in mind that we have independent reasons for being suspicious about the term “temporal instant *t*”. Consequently, in the process of CA, we should go on to define the relationship between parts and whole as follows: over some period of time, the Xs jointly compose Y, and over that period of time, none of the Xs completely overlap Y.

With respect to the notion that either Y or X (or both) exists *right now*, the standard construal of composition, which implies that “now” or “right now” is to be understood in terms of ontological synchronicity is fundamentally ill-equipped for the analysis of complex and dynamical phenomena such as the example of returning a tennis serve and the temporal trajectory of an action potential. That is, nothing in the standard account of composition allows for *continuous processes unfolding in real time*. And both returning a tennis serve and the manifestation of an action potential *are* processes that unfold continuously (over a certain time-course) in real time. Clark provides the following definition of a continuous process:

“A continuous process is one in which the time-series of explanatorily relevant sub-states cannot be reduced to a sequence of discrete states with jumps in between, but instead requires a genuine continuum of states.” (1998, 356)

In the tennis example, it makes little sense to insist that Y is composed *wholly within and only within* each particular sub-interval of *t* and that each transition from one sub-interval to the next involves a complete transition of X and Y such that both X and Y are *wholly present* within one and only one particular sub-interval *at a time*⁵.

Similarly in the debate about CA, I submit. For example, when Clark adopts the short-term timescales of the here-and-now to explore just how the brain participates in what (from the perspective of EC) are new distributed cognitive products, Clark wants to analyze which neural and bodily processes assemble temporally distributed wholes *right now*. Indeed, as Clark mentions: “In depicting the processes of on-the-spot recruitment and exploitation as neurally-centered, I meant only to stress the pivotal role, on all these shorter time-scales, of the specifically neural changes that immersion in those cultural practices presumably inculcate.” (2011, p. 459) For instance, in *Supersizing the Mind* (2008), most of the case studies referred to by Clark are studies that emphasize the short-term, but varied, temporal scales of bodily and neural operations – timescales that unfolds over courses of 50 to 300 milliseconds.

In contrast to the standard view of composition in metaphysics, when Clark states that it is the processes, which operate in the here-and-now that assemble (or, compose) distributed ensembles, what Clark is actually saying is that it is the short-term timescales over which most neural and bodily processes operate that *during that short period of time* assembled or put together some distributed cognitive whole.

Here it is enlightening to consider that the use of “right now” in the process of CA is closely related to how the folk consider the notion “right now”. In EC, occurrent distributed cognitive wholes are considered to temporally unfold everywhere from 50 milliseconds and up to a few hours in the case of occurrent emotions and extended instances of decision-making. Similarly, and as I argued above, if the folk are right, then the term “right now” would refer to occurrences ranging from 2-3 minutes in their entirety to 200-300 milliseconds. That is, from processes involved in complex decision making to the completion of one saccadic eye movement. Furthermore, as with most (if not all) processes, cognitive processes have

⁵ As Spivey mentions: “Real time does not function like a digital computer’s clock. It does not move forward and then stop to be counted, and then move forward again only to stop again. At the level of human behavior, real time does not have an objective functional unit.” (2007, p. 30)

subprocesses as well as subphases. For example, the typical time span of a single episode of voluntary biographical remembering is roughly 10 seconds and this trajectory can have any number of both continuous and discontinuous sub-phases and sub-processes. Thus, just as the processes making up a tennis serve return, the processes and subprocesses that compose a token episode of voluntary biographical remembering are *time continuous* – that is, both cases involve continuous processes, and continuous processes and their relevant subprocesses as well as subphases cannot be reduced to a sequence of discrete states instantiated within a discrete temporal slice of time. This idea is nicely expressed by Spivey, who says:

“[Claiming] that a system was in a particular “state,” **X**, at a particular point in time, really boils down to saying that the *average* of the system’s states during that *period of time* was **X**. This kind of coarse averaging measurement is often a practical necessity in science, but should not be mistaken as genuine evidence for the system actually resting in a discrete stable state.” (2007, p. 30; bold and italics in original)

It is, of course, possible to distinguish conceptually between long-term evolutionary timescales, time-scales running over developmental and/or cultural-historical time, and the short-term timescales of hours, seconds, and milliseconds (see e.g., Clark 2011)⁶. But, regardless of how we conceptually carve up time, time is, I submit, continuous (Clark 1998; Port & van Gelder 1995; Spivey 2007). Indeed, it seems highly unlikely that the components assembled on the spot to complete a tennis serve return function in what van Gelder & Port call *arbitrary step time* (t1, t2, etc.). That each new second or millisecond (or smaller) signaled, as Spivey puts it “an instantaneous *and simultaneous* updating of the discrete state of each and every unit

⁶ But these conceptual distinctions may not apply objectively. As Smart (1963), for instance, argues against the A-theory of time, according to which “past”, “present” and “future” are understood to objectively apply to the universe, Smart argues that this way of carving up time is an entirely anthropocentric account of time. That is, distinctions of past, present, and future are distinctions made from a particular (human) point of view (Smart 1963, p. 132; see also Sider 2001; for an overview of the A-theory of time, see e.g., Mellor 1998)

in the system.” (2007, p. 30; italics in original) This complaint applies to the standard view of synchronic composition, in that, this standard view delineates time into discrete chunks with jumps in between them. In the standard cases, the Xs compose or assemble Y only if the Xs and Y are wholly present at each particular instant in time at which they exist (*cf.* van Inwagen 1990).

The problem, in short, is that even though both the SCQ and the process of CA address the question ‘under which circumstances, and by which principles’ do certain entities compose or assemble other entities, the SCQ is formulated synchronically and, therefore, leaves out the temporal dynamics of actually occurring instances of composition. What we really need is a temporally quantified version of the SCQ; call it the *Temporal Special Composition Question* (TSCQ). Unlike the SCQ, the TSCQ does not presuppose a temporal restriction on the verb “compose” or on the verb “assemble” such that these would imply “compose or assemble at this very instant”. Instead, the TSCQ asks the question “over which timescales do processes operate when they jointly compose (or, assemble) a whole?”

In addition, “wholes” may be temporary and a one-off ensemble, a temporary and repeatable one, or something more permanent (*cf.* Wilson & Clark 2009). Whether it is one or the other is an empirical question – not a question to be settled by metaphysical analysis.

If I am right that the claim “the Xs compose or assemble Y right now”, at least when considering time continuous processes, is an abstraction and should not be mistaken as evidence for the claim that some Xs assemble Y at an ontologically synchronic instant, then the evaluation of empirical evidence supporting my claim must be sensitive to this fact. Fortunately, the empirical evidence is sensitive to this fact. Here, then, is another difference between the understanding of “time” in the standard account of composition, on the one hand, and the kind of composition we find in time continuous systems. In the former, time is portrayed as a dimension that is *neutral*, i.e., time exists independently of the events or states, or processes, etc., that occur in time, while time plays a *fundamental role* in the latter. For example, in their discussion of the role of the body (or, embodiment) in cognition, Ballard and colleagues say the following:

“When the production of intelligent behavior by the body-brain system is taken into account, the constraints of time and space intervene to limit what is possible.” (Ballard et al. 1997, p. 723)

If I am correct, the synchronic composition view treats time only as a specification of the proposition “the Xs compose Y” as taking place at a time. Thus, the locution “P is v at t ” implies that the expression in place of P refers to a proposition (e.g., the Xs compose Y), the expression in place of v refers to a truth value (it is either true or false that P), and the expression in place of t refers to a particular time instant such that on the standard view “P is v at t ”.

To proceed further with this temporality-driven critique of synchronically formulated composition, and to hook it up with the discussion of CA, what we need, in the context of CA, is positive empirical evidence that time really matters for just how and for which processes are assembled in order to solve a given problem⁷. That is, we shall look at an example that Clark argues shows the “balanced use of a set of potentially highly heterogeneous resources assembled on the spot to solve a given problem.” (2008, p. 13)

5. Cognitive assembly over short-term temporal frequencies

Consider, then, Clark’s employment of a series of experiments conducted by Gray & Fu (2004) targeting how patterns of interactive behavior emerges at the level of embodiment and how *soft-constraints* – at the embodiment level – determine which of the possible strategies – for solving a given problem – are most likely to be selected given the problem or task environment.

⁷ One might object to my claim that dynamical cognitive science is incompatible with tenseless accounts of time, in that, you can account for change in tenseless terms as Russell famously showed. Briefly, what it is for an entity E to undergo change is for E to have a property X at t and a property Y at t_1 rather than X at t_1 . But, notice, if we want to understand (i) the evolvment of the system from t to t_1 , and (b) how that particular temporal evolvment gives rise to a property difference in E from t to t_1 , then a synchronic explanation comes up short.

A few points of clarification: as other researchers in embodied cognition (see e.g., Ballard et al. 1997), Gray & Fu take the notion “embodiment level” to refer to the timescales over which several neural and bodily operations begin to cohere into certain patterns of activity that compose or assemble the bases of interactive behavior. These operations include what Gray & Fu refer to as elementary cognitive, perceptual, and action operations that have a typical time-course of 300 milliseconds (2004, p. 362) Thus, in this case, and with respect to the TSCQ, the timescales over which processes operate when they jointly compose or assemble some short-term cognitive products to solve a given problem is the short-term timescale of 300 milliseconds it takes various neural processes to combine with each other to form a specific “microstrategy” (cf. Gray & Fu 2004, p. 364) – where “microstrategy” refers to patterns of behavior invoked to accomplish a cognitive task. Such microstrategies, Gray & Fu stress, are *softly constrained*, suggesting that there are many possible routes rather than just one (hard constrained option) by which various neural/bodily processes may combine or come together in order to solve a given cognitive task (Gray & Fu 2004, p. 361).

In the first set of experiments (Gray & Fu 2004), subjects were presented with the task of having to program an on-screen simulation of a VCR control panel. The idea of the experiment was to manipulate the time-course and time-cost involved in accessing the information required to program a VCR in order to assess whether the task environment facilitates or discourages the use of “knowledge in-the-world for knowledge in-the-head.” (Gray & Fu 2004, p. 364)

In the experiment subjects were divided into three groups. In the first group (the Free-Access condition), the information was clearly visible in front of the user so that she freely could access the information via saccadic eye movement. In the second group (the Gray-Box condition), the window was partly visible, although the required information (about channel, start time, etc.) was covered with a gray box. To uncover the information the user had to remove the gray box via a mouse click on the gray box. The final group (the Memory-Test condition), who, unlike the others, had previously memorized all the information required, had to remove the gray box and type in the necessary information. In order to determine to time-course and time-cost involved in each of these three conditions, Gray & Fu analyzed two components: first, the time needed for *perceptual-motor access* to the information; and second, the time needed for *memory retrieval* (see Fig. 1). What Gray & Fu found was that time costs

of information retrieval, measured in milliseconds, are what determine the combination of processes (biological memory, motor actions, shifts of attention, etc.) assembled to solve the problem. As Clark puts it:

“[The] subjects settled on whatever strategy yielded (at that phase of the programming) the least cost (measured by time) information retrieval. In fact, they did this even when the fastest mix of resources sacrificed perfect knowledge in the world for imperfect knowledge in the head. Only when the in-the-world data could be accessed with less effort (measured by time) than the data stored in biological memory was it recruited and were calls to the external store “built into” the dominant strategy.” (2008, p. 119)

Estimates (in ms) of perceptual-motor and memory retrieval effort by condition		
Condition	Perceptual-motor access	Memory retrieval
Free-Access	500	500-1,000
Gray-Box	1,000-1,500	500-1,000
Memory-Test	1,000-1,500	100-300

Fig. 1 Overview of the time needed to access the information on each condition. Estimates are in milliseconds (adapted from Gray & Fu 2004, p. 368).

If Gray & Fu are correct, what this example clearly indicates is that psychological phenomena of using softly constrained patterns of information retrieval cannot be appropriately explained without (a) an appeal to richly temporal processes, and (b) insofar as various cognitive, perceptual, and motor elements combine to jointly compose such transient microstrategies. Thus, the timescales over which the Xs operate when they jointly compose Y fail to accommodate the temporally restricted assumption that the Xs do so at *right this instant t*. Indeed, rendering the verb “compose” in “the Xs compose Y” either tenseless (as the non-presentist would insist on) or in present tense (as the presentist would insist on) makes little sense in this dynamical and time continuous domain.

First, as Clark mentions, temporal “cost-benefit trade-offs are said to provide a soft constraint [...] on the mix of motoric, perceptual, and biomemory-based resources that will, other things being equal, be automatically recruited to perform a given information-processing task on a given occasion.” (2008, p. 120) That is, Gray & Fu show that the Free-Access group favor perceptual-motor access over memory retrieval, whereas the Memory-Test group favor memory retrieval strategies given that that route of retrieval is much faster than perceptual-motor access. It would seem, then, that the timescales over which various cognitive, perceptual, and motor elements combine to jointly compose transient microstrategies are ineliminably *context-sensitive* and will consequently differ accordingly to the constraints of the task environment. That is, it would seem to be impossible to require that each case of the verb “compose” in the expression “the Xs compose Y” must be true or false once and for all, that is, independently of time and context.

Second, processes preclude instantiation at a particular time instant t , in the sense that what is it to *be* a process – that is, what it is to persist as a process – involves, necessarily, unfolding over time (*cf.* Hofweber & Velleman 2011). Therefore, even on the short-term timescales of neural operations – note also our discussion of the action potential earlier – neural processes cannot be completely and wholly present at any single instant. That is because neural processes, like processes in general, are temporally extended in nature. For my purposes, then, the real power of the example discussed by Clark is that it shows that even on the short-term timescales over which neural and bodily processes operate, it is ontologically diachronic *all the way down*.

We have here a consequence for those involved in the debate over the process of CA in EC and for metaphysicians with a synchronic persuasion. First, and to repeat what I said in section 2, insofar as Clark states that it is the processes that operate *here-and-now* that orchestrate the assembly of hybrid, distributed cognitive wholes, this claim really boils down to saying that it is the short-term timescale of bodily and neural processes that *during that period of time* orchestrated the assembly process of some distributed cognitive whole, and *should not be mistaken as genuine evidence for the processes actually assembling some distributed cognitive whole at a time instant t* . As I understand Clark’s position – or, the best way to interpret Clark’s insistence on the timescale of the here-and-now – is precisely that it is the processes that *unfold over* short-term, but varied, timescales that assemble or compose distributed cognitive processes and/or systems (Clark 2011). Second, if the metaphysics of composition is

to apply to dynamical processes involved in the process of CA, then the synchronic account of composition is ill equipped. Thus, the process of CA must be stated entirely such that it does not implicate assumptions about composition as ontologically synchronic.

6. Cognitive assembly – distributed over multiple timescales and multiple resources

So far I have argued that any adequate analysis of the TSCQ in the context of CA should accept a diachronic conception of composition. I wish now to examine what is presupposed in the argument for the process of CA provided by Clark, namely that it is only the processes operating here-and-now that are responsible for the assembly of distributed cognitive processes and/or systems in conjunction with the claim that those processes responsible for such assembly are bodily and neural processes.

As Clark explicitly states: “My own targets, in the discussion in [Supersizing the Mind] of cognitive assembly, were the processes operating in the here-and-now.” (2011, p. 459) And as Clark specifies the kinds of processes in operation in the here-and-now:

“It is indeed primarily (though not solely) the biological organism that, courtesy especially of its potent neural apparatus, spins and maintains (or more minimally selects and exploits) the webs of additional structure that then form parts of the machinery that accomplishes its own cognizing.” (2008, p. 123)

There are two assumptions at work in Clark’s project. The first assumption is that only processes operating here-and-now (or right now) are responsible for *the assembly* of distributed cognitive processes or systems. The second assumption is that the processes most directly responsible for such assembly are bodily and/or neural processes. With respect to the example above, these two assumptions highlight that it is the short-term temporal frequencies at the embodiment level, made up of perceptual, motor, and cognitive processes, which primarily compose higher-level cognitive products.

However, this combination is not the only coherent and live option in the literature. That is, if we call Clark’s first assumption for “A”, and call the second

assumption “B”, then these assumptions leave open any of the three coherent possible combinations: (i) accept (A) and (B) – this is Clark’s position in his (2008); (ii) accept (A) but deny (B); and (iii) deny both (A) and (B). That is:

1. Only processes operating here-and-now are responsible for CA, and such processes are bodily and neural processes.
2. Only processes operating here-and-now are responsible for CA, but this does not prevent non-neural and non-bodily processes from significantly contributing to the process of CA.
3. Neither processes operating here-and-now nor bodily and neural processes are primarily responsible for CA.

All three of these combinations have seen defenders in the contemporary literature, although I suspect that whether or not there really are exponents of the third combination depends on interpretations of key selected passages. For example, Clark (2011) has attacked Hutchins’ “Hypothesis of Enculturated Cognition” (Hutchins 2011) for presupposing that “cultural practices are *sufficient* to account for all the crucial work of cognitive assembly.” (2011, p. 459; italics in original) And as Clark specifies: “I think Hutchins is failing to attend to important differences concerning the shape and timescale of the processes concerned.” (2011, p. 459) By stating that Hutchins is failing to attend to different aspects concerning timescales, Clark means:

“Hutchins response might be that we should simply reject the conceptual separation between the processes operating on [...] various timescales. [...]. That is how I read his key suggestion that “both the constraints of cultural practices and the malleable internal microdemons can be seen as elements of a single adaptive system”. But while I agree that these are indeed (also) elements of a single long-term adaptive system, that does nothing to diminish the conceptual separation between the long-term evolution of cultural practices, the medium-term effects of my immersion in such practices, and the short-term processes by means of which my brain then participates in what (from an

extended mind perspective) are new hybrid cognitive routines that productively criss-cross brain, body, and world.” (2011, p. 460)

On this interpretation of Hutchins’ position, it appears that Hutchins endorses the third combination, namely that neither processes operating in the here-and-now nor bodily and neural processes are what primarily assembles distributed cognitive wholes. Hutchins himself states that one way to avoid the option of combining (A) and (B) is “to abandon the assumption that the biological brain is the essential element. Doing so, of course, requires that one look elsewhere for the apparently impartial forces that assemble cognitive systems.” (2011, p. 439) As Hutchins proposes: “A good start to understanding this process of recruitment would be to notice the role of cultural practices in the orchestration of soft-assembly of extended cognitive systems.” (2011, p. 440)

A worry about the third combination, however, is that the general formulation of this combo allows for the following claim: that certain evolutionary conditions for some present cognitive functioning (e.g., from 100,000 years ago) could be actively orchestrating the process of CA here-and-now. I am not aware of any philosophers that have defended such a suggestion but it is within the logical possibilities of the general formulation of the third option.

However, in other passages, Hutchins’ own position is much closer to the second option. Clark mentions this possibility as well, as he says: “For as Hutchins himself says, it is only the ‘special super-flexible medium’ of the brain that *allows* such shared practices to come to orchestrate human learning and response in the first place.” (2011, p. 459; italics in original) Or, as Hutchins states:

“In this perspective, the brain appears as a special super-flexible medium that can form functional subsystems that establish and maintain dynamic coordination among constraints imposed by the world of cultural activity, by the body, and by the brain’s own prior organization.” (2011, p. 445)

Thus, depending on how one interprets Hutchins’ position, it is possible to place him in either option two or three. Now, even though I say that all three combinations are coherent, it should be clear that any one of these is not free of difficulty. In addition to Clark’s criticism of Hutchins’ hypothesis of enculturated cognition, Hutchins has

argued that the first combination – the one Clark opts for in much of his latest work (Clark 2008, 2011) – is problematic, in the sense that endorsing the view that only bodily and neural processes are responsible for the process of CA excludes from view that much of the “heavy-lifting” – as Hutchins is fond of calling it – in the assembly of distributed cognitive ecologies is performed by cultural practices that unfold over longer timescales than those of the here-and-now.

A similar sort of ambiguity is present in Clark’s authorship, especially when one compares some of Clark’s earlier work such as *Being There* (1997) as well as articles such as “Beyond the Flesh: Lessons from a Mole Cricket” (2005a) and “Word, Niche and Super-Niche: How Language Makes Minds Matter More” (2005b) and his latest book *Supersizing the Mind* (2008). In fact, one may also locate a similar kind of ambiguity just by reading through his (2008). The ambiguity consists in the following: whereas Clark in his earlier works is much closer to the combination “accept (A) but deny (B)”, he is much more in favor of, as we have seen, the combination “accept (A) and (B)” in his latest work. Consider, for example, the following combination of quotes: first, “[my] own targets, in the discussion in [Supersizing the Mind] of cognitive assembly, were the processes operating in the here-and-now.” (2011, p. 459) And as Clark specifies where those processes are primarily located: “It is indeed primarily (though not solely) the biological organism that, courtesy especially of its potent neural apparatus, spins and maintains (or more minimally selects and exploits) the webs of additional structure that then form parts of the machinery that accomplishes its own cognizing” (2008, p. 123) So, here Clark is explicit about endorsing option one – that is, the acceptance of (A) and (B).

But, consider, then, the following couple of quotes, the first from his (2008) and the second from his (1997): first, “[this] is not to deny, of course, that much of the spinning is done by social groups of organisms spread out over long swaths of history.” (2008, p. 243; footnote 18) And second: “[The] brain in its bodily context, interacting with a complex world of physical and social structures. These external structures both constrain and augment problem-solving activities of the basic brain, whose role is largely to support a succession of iterated, local, pattern-completing responses.” (1997, p. 191) Thus, whereas the first set of quotes puts Clark squarely in the first option, this latter set of quotes puts him firmly in the second option.

It is certainly true that one way to read these ambiguities in both Clark’s and Hutchins’ work is that they indicate that not any one of the three options can be

defended on metaphysical ground. That is, it is an empirical question just how often and how much of the assembly work is performed through the combination of (A) and (B) or (A) and not (B). It is an empirical issue how much and how often the integration is orchestrated internally and how much and how often the integration is assembled externally. This, I suspect, is the correct way by which to understand the ambiguity present in both Clark and Hutchins concerning the process of CA.

However, this raises a different question, namely might there be reasons for favoring one option or combination over the other even though the two combinations are not mutually exclusive *per se*? I think that there are such reasons, especially reasons that lend support to the combination “accept (A) but deny (B)”. One worry with the first combination is that it threatens to screen-off the fact that even in the here-and-now, history and culture are *always already* embedded as well as carried along in the practices and artifacts individuals are engaging with (Menary 2007; Sutton 2008, 2010; see also Haugeland 2002; Lave & Wegner 1991). Indeed, the option of taking the second combination on-board is much more in line with a distinctive third-wave of EC theorizing (*cf.* Sutton 2010; Cash 2013; Kirchhoff 2012). That is, even in the here-and-now, across the short-term timescales of hours, minutes, seconds, and so on, the process of EC is not primarily driven/orchestrated by bodily or neural processes but it also significantly morphed and sculptured by socially embedded and culturally transmitted practices.

As I have shown, Clark’s position is (in certain works) consistent with the second option, thereby bringing Clark into contact with a third-wave of EC research. Because of this, I shall consider a case study discussed by Clark on how expert bartenders, when faced with a multiple drink order in noisy and crowded environments are able to successfully solve the problem⁸. Since Clark articulates this example within the framework of niche construction (NC) (*cf.* Laland et al. 2000), I start by giving a brief introduction to the central tenets of NC. NC, as defined by Laland et al. (2000), refers to:

“[The] activities, choices and metabolic processes of organisms, through which they define, choose, modify and partly create their own niches. [...]. For example, to varying degrees, organisms choose their own habitats, mates, and

⁸ The actual case study is due to Beach (1988).

resources and construct important components of their local environments such as nests, holes, burrows, paths, webs, dams, and chemical environments.” (2000, pp. 132-33)

Organisms adapt to environmental pressures. But organisms also construct, alter, and modify their own environmental niches. Some make burrows, webs, shelters, and other resources. Earthworms are a good example of what we might call *pragmatic engineering*. Earthworms engage in burrowing activities, often resulting in a transformation of the structure and chemistry of the soil in which they live (Laland et al. 2000, p. 134). This burrowing activity is important because earthworms, prior to their presence on land, were originally aquatic organisms (Laland 2004, p. 321). As Laland says: only by “co-opting the soils that they inhabit and the tunnels they build to serve as accessory kidneys that compensate for their poor structural adaptation [...],” (2004, p. 321) can earthworms tackle the physiological demands of a different water- and salt-balance on land.

Many organisms not only alter and transform their environmental niches pragmatically. Much niche construction is importantly a mode of *epistemic engineering*, in the sense that active niche-constructors modify and alter the informational character of the environment (Sterelny 2010, p. 470). For instance, ants lay scent trails between nest and food source. Humans off-load information “onto” the environment so as to ease the burdens on “internal” memory processing (Donald 1991). Other organisms, like hawks, simply choose the best spot from which to maximize the view of their hunting territory. In a comprehensive study on the “intelligent use of space,” Kirsh argues that rearranging spatial relations between environmental resources transforms the problem solving space by reducing the descriptive complexity of the task environment (1995, 2009). Especially in the human lineage, the ramification of epistemic engineering is the establishment of a cumulatively constructed cognitive-developmental niche (Sterelny 2003, 2010).

Epistemic (and pragmatic) engineering is not only cognition-enhancing in the heat of some problem-solving scenario, since some modifications to the physical and informational environment are transmitted downstream to the following generation. As Sterelny argues, “cumulative downstream epistemic engineering” implies transmission of both socio-cultural structures and ecological and technical know-how or expertise enabling the transmission and acquisition of new knowledge (2010, p.

470). In all these cases of NC, both pragmatic and epistemic varieties, what matters, as Laland et al. (2000) emphasize, is that the activity of NC leads to new *feedback cycles*.

In the standard cases of NC, those feedback cycles run over evolutionary timescales. However, for Clark's purposes, what really matters is that "this whole process has a direct analogue within lifetime learning." (2005a, p. 256) As Clark states:

"Here, the feedback cycles alter and transform the processes of individual and cultural learning. For example, both educational practices and human-built structures (and artifacts) are passed on from generation to generation in ways that dramatically alter the fitness landscape for individual lifetime learning." (2005a, p. 256)

The example that Clark considers is how an expert bartender, when faced with a multiple drink order in very noisy and crowded circumstances, is capable of solving the problem successfully. Or, alternatively, just how learning within a pre-structured niche with cultural practices and differently shaped glasses, makes it possible for a novice bartender to perform competently. In lifetime learning, or across developmental timescales, the expert bartenders learn how to line up differently shaped glasses in spatial sequences, which, in turn, correspond to the temporal sequence of drinks orders. As Clark states: "The problem of remembering what drink to prepare next is thus transformed, as a result of learning within this *pre-structured niche*, into the problem of perceiving different shapes and associating each shape with a kind of drink." (2005a, p. 256; italics added) In this sense, the cultural practices of knowing what to do, and how to do it, when facing a multiple drink order are shaped by the niche constructing activity of previous individuals, and these practices constrain our epistemic access to the world by orchestrating what to attend to and see when so attending (*cf.* Hutchins 2008). This resonates deeply with Haugeland's idea that normative practices have a certain kind of "normative gravity" (2002, p. 32). In becoming normalized in the practice of bartending, if we take Haugeland's view, what is normalized is not so much behaviors but rather *dispositions to behave*. But even if normal practices or behaviors might never be exactly alike, they are sufficiently alike to be within the same "orbit". Thus, according to Haugeland: "when

an individual's dispositions stray from producing behavior within these orbits (that is, types [of normative practices]), they are 'pulled back in'." (2002, p. 32) Another way of articulating this idea is due to the patterned practice approach developed by Roepstorff et al (2010). As they mention: "From the inside of a [cultural] practice, certain models of expectancy come to be established, and the patterns, which over time emerge from these practices, guide perception as well as action." (2010, p. 1056)

Here it is not primarily bodily and neural processes that compose a bartender's ability to get the job done but rather neural processes in coordination with normative, cultural practices that jointly come together to assemble such abilities.

7. Implications

The combination of ontological diachronicity with the TSCQ implies that whenever the statement "the processes involved in the process of CA are those processes operating in the here-and-now" is made, it follows that we must understand this statement as expressing "the processes involved in the process of CA are those processes operating in the here-and-now *during that period of time*". This particular view, it seems to me, is the implicit view of several philosophers of cognitive science as well as some cognitive scientists themselves (see e.g., Ballard et al. 1997; Beer 2000; Chemero 2009; Clark 1998b; Kirchhoff 2013; Spivey 2007; Varela et al. 2001; and others). One important implication this has for any synchronic notion of composition – such as the SCQ – is that not only is it ill fitted to analyze temporally complex phenomena; it can never be made to analyze such temporal phenomena *simpliciter*.

Once we make room for a robust diachronic account of the circumstances under which entities of different kinds assemble or compose another entity, and once room has been made for the pivotal role of cultural practices in this process of assembly or composition, this requires (non-trivially) that we look beyond the system made up of the individual agent and artifact. Notice that there is nothing special about endorsing option 1 above: only processes operating in the here-and-now are responsible for the process of CA and such processes are primarily bodily and neural processes. Indeed, prominent advocates of this view include Clark (2008, 2011) and Clark & Chalmers (1998). As we saw in introduction, this particular version of EC is what both Menary (2010) and Sutton (2010) refer to as *first-wave EC*. Defenders of first-wave EC

ground their arguments for EC on the ideas that (i) external artifacts are recruited into the cognitive systems of an individual due to the right kind of causal coupling, and (ii) functional similarity between the causal roles of internal and external physical occupiers. If the arguments for EC focuses on how artifacts are integrated into an individual's cognitive system, then it is not surprising that first-wave versions of EC usually adopt an account of the process of CA along the lines of option 1. That is, in Hutchins' words, "if culture is reduced to a collection of lifeless artifacts" (2011a, p. 444), then the active dynamic processes involved in the process of CA must be bodily and neural. But, if we cultural practices may be conceived of as playing a central role in the process of CA, then it follows that some of the active dynamic processes involved in the process of CA lies beyond the system made up of the individual agent and artifact. If this turns out to be correct, then what we have is an account of the process of CA that grounds Sutton's gesture towards a third-wave version of EC: a version "which dissolves individuals into peculiar loci of coordination and coalescence among multiple structured media [and practices]". (2010, p. 213)

I wish to finish this paper by considering whether it is possible to apply this metaphysical contribution to the debate over the process of CA to settle any disputes in the literature. I think that this is entirely possible. Consider, for example, how Hutchins attempts to push Clark into a strictly neural-oriented position with regards to the process of CA by exploiting an apparent bias in the phrase "on the spot". As Hutchins says:

"According to Clark, this exploitation happens "on the spot," but the constraints that determine which resources are exploited and how they are related to one another is not entirely formed "on the spot". The "on the spot" phrase highlights the opportunistic nature of cognitive systems. However, without additional discussion, this wording may also bias the solution toward the biological brain by isolating the activity from the context of cultural historical processes." (2011, p. 441)

One cause for concern about Hutchins' interpretation of the phrase "on the spot" is that there is nothing about the phrase "on the spot" that conceptually entails a commitment to the view that it is the brain that is the most active element in the

assembly of distributed cognitive products. An example will make this more concrete. Consider, again, a passage from Hutchins:

Cultural practices shape active sensing and ways of seeing the world by highlighting what to attend to and what to see when so attending. Clark mentions the activity of seeing a star. A far more interesting example is seeing a constellation, since a constellation exists only by virtue of someone enacting it via a cultural practice that allocates visual attention in a particular way.” (2011, p. 441)

Whichever processes combine to produce the capacity to see a star constellation do so *on the spot* – on the timescales of seconds or, perhaps, minutes. The question is: over which other timescales would such processes be active?

Instead of juxtaposing the short-term timescales and long-term timescales (e.g., historical time-scales), Hutchins would be better off arguing that there is no problem with depicting the process of CA as unfolding over the short-term timescales of the *here-and-now*, provided that you leave room for the central roles of cultural practices (or, history embedded in those practices) *in* the processes that unfold here-and-now. Insofar as the meaning of “now” is such that it may, in the right circumstances, include a dynamical interval of time, and insofar as the cultural practices within which the cognitive task is carried out unfolds within such a dynamical interval of time, then cultural practices may be part of the processes assembling some cognitive ability. That is the real point that one will be able to make by opting for combination (A) and not (B) above. Consequently, Hutchins cannot appeal to the notion “on the spot” in order to assert that Clark privileges the brain in the process of CA, since the meaning of “on the spot” is contingent – as I have argued – on the time and place of the utterance.

8. Conclusion

What I hope to have shown in this paper is that when considering the process of CA, that is, when we analyze over which timescales certain processes combine to compose distributed cognitive products, there is no ontologically synchronic instant t at which that is possible. I wish to finish this paper by pointing out that metaphysical analysis

cannot settle the question over which timescales the processes involved in the process of CA are predominantly active. This is a matter of empirical investigation. However, by scrutinizing the metaphysics of what it means for certain Xs to compose a certain Y now, it is possible, I think, to turn what might look like a metaphysical dispute into a productive recipe for empirical research and to set certain constraints for how such research must be carried out.

References

- Ballard, D., Hayhoe, M., Pook, P., & Rao, R. (1997). Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences*, 7(1), 66-80.
- Beach, L. (1988). The role of external mnemonic symbols in acquiring an occupation. In M. M. Gruneberg and R. N. Sykes (eds.), *Practical Aspects of Memory*, vol. 1. New York: Wiley.
- Bechtel, W. (2008). *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. New York & London: Routledge
- Beer, R (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72, 173-215.
- Bennett, K. (forthcoming). *Making Things Up*. Under contract with *Oxford University Press*.
- Bennett, K. (2011). Construction area (no hard hat required). *Philosophical Studies*, 154, 79-104.
- Cash, M. (2013). Cognition without borders: “Third-Wave” socially distributed cognition and relational autonomy. *Cognitive Systems Research*: DOI: <http://dx.doi.org/10.1016/j.cogsys.2013.03.007>
- Chemero, A. (2009). *Radical Embodied Cognitive Science*. Cambridge, MA: The MIT Press.
- Clark, A. (2011). Finding the Mind. *Philosophical Studies*, 152, 447-461.
- Clark, A. (2008). *Supersizing the Mind*. Oxford University Press.
- Clark, A. (2005a). Word, niche and super-niche: How language makes minds matter more. *Theoria*, 54, 255-268.
- Clark, A. (2005b). Beyond the flesh: Some lessons from a mole cricket. *Artificial Life*, 11, 233-244.
- Clark, A. (1998). Time and mind. *Journal of Philosophy*, XCV(7), 354-376.

- Clark, A. (1997). *Being There*. Cambridge, MA: The MIT Press.
- Craver, C.F. (2007). *Explaining the Brain: mechanisms and the mosaic unity of neuroscience*. Oxford: Clarendon Press.
- Donald, M. (1991). *Origins of the Modern Mind*. Cambridge, MA: Harvard University Press.
- Engel, A., Pascal, F., & Singer, W. (2001). Dynamic predictions: Oscillations and synchrony on top-down processing. *Nature Reviews*, 2, 704-716.
- Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Boston: Houghton-Mifflin.
- Gray, W.D., and W-T. Fu. (2004). Soft constraints in interactive behavior: The case of ignoring perfect knowledge in the world for imperfect knowledge in the head. *Cognitive Science*, 28(3), 359-382.
- Heller, M. (1984). Temporal parts of four-dimensionalism. *Philosophical Studies*, 46, 323-334.
- Haugeland, J. (2002). Andy Clark on cognition and representation. In H. Clapin (ed.), *Philosophy of Mental Representation* (pp. 24-36). Oxford, UK: Oxford University Press.
- Hodgkin, A. L., & Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*, 117, 500-544.
- Hutchins, E. (2011). Enculturating the Supersized Mind. *Philosophical Studies*, 152, 437-446).
- Hofweber, T., & Velleman, D. (2011). How to endure. *The Philosophical Quarterly*, 61(242), 37-57.
- Kirchhoff, M. D. (2013). Extended cognition & the causal-constitutive fallacy. In search for a diachronic and dynamical conception of constitution. *Philosophy and Phenomenological Research*, 1-41. DOI: 10.1111/phpr.12039
- Kirchhoff, M. D. (2012). Extended cognition and fixed-properties: Steps to a third-wave version of extended cognition. *Phenomenology and the Cognitive Sciences*, 11(2), 287-308.
- Kirsh, D. (2009). Problem solving and situated cognition. In P. Robbins & M. Aydede (eds.), *The Cambridge Handbook of Situated Cognition* (pp. 264-306). Cambridge: Cambridge University Press.
- Kirsh, D. (1995). The intelligent use of space. *Artificial Intelligence*, 73, 31-68.

- Laland, K. N. (2004). Extending the Extended Phenotype. *Biology and Philosophy*, 19, 313-325.
- Laland, K. N., Odling-Smee, J., & Feldman, M. W. (2000). Niche construction, biological evolution, and cultural change. *Behavioral and Brain Sciences*, 23, 131-175.
- Lave, J., & Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*. Cambridge, UK: Cambridge University Press.
- Lewis, D. (1986). *On the Plurality of Worlds*. Oxford: Basil Blackwell.
- Machamer, P., Darden, L., and Craver, C.F. (2000). Thinking about mechanisms. *Philosophy of Science*, 67(1), 1-25.
- Markosian, N. (2004). A defense of presentism. In D. W. Zimmerman (ed.), *Oxford Studies in Metaphysics*, vol. 1 (pp. 47-82). Oxford: Oxford University Press.
- Markosian, N. (1998). Brutal composition. *Philosophical Studies*, 92, 211-249.
- Mellor, D. H. (1998). *Real time II*. London and New York: Routledge.
- Menary, R. (2007). *Cognitive Integration: Mind and Cognition Unbounded*. Basingstoke: Palgrave Macmillan.
- Merricks, T. (2001). *Objects and Persons*. Oxford: Clarendon Press.
- Olson, E. T. (2006). Temporal parts and timeless parthood. *Noûs*, 40(4), 738-752.
- Roepstorff, A., Niewöhner, J., & Beck, S. (2010). Enculturing brains through patterned practices. *Neural Networks*, 23, 1051-1059.
- Rosen, G., & Dorr, C. (2002). Composition as a fiction. In G. Richard M (ed.), *The Blackwell Guide to Metaphysics* (pp. 151-174). Oxford: Blackwell.
- Rupert, R. (2009). *Cognitive Systems and the Extended Mind*. Oxford: Oxford University Press.
- Sider, T. (2001). *Four-dimensionalism*. Oxford: Oxford University Press.
- Smart, J. J. C. (1963). *Philosophy and Scientific Realism*. London: Routledge & Kegan Paul.
- Spivey, M. (2007). *The Continuity of Mind*. Oxford and New York: Oxford University Press.
- Sterelny, K. (2010). Minds: Extended or Scaffolded? *Phenomenology and the Cognitive Sciences*, 9(4), 465-481.
- Sterelny, K. (2003). *Thought in a Hostile World*. Malden, MA: Blackwell Publishing.

- Sutton, J. (2010). Exograms and interdisciplinarity: history, the extended mind, and the civilizing process. In R. Menary (ed.), *The Extended Mind* (pp. 189-225). Cambridge, MA: The MIT Press
- Sutton, J. (2008). Material agency, skills and history: Distributed cognition and the archaeology of memory. In C. Knappett and L. Malafouris (eds.), *Material Agency* (pp. 37-55). Springer.
- Van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, 21, 615-665.
- Van Gelder, T., & Port, R. (1995). It's about time: An overview of the dynamical approach to cognition. In R. Port and T. van Gelder (eds.), *Mind as Motion: Explorations in the Dynamics of Cognition* (pp. 1-44). Cambridge, MA: The MIT Press.
- Van Inwagen, P. (1994). Composition as identity. *Philosophical Perspectives*, 8, 204-220.
- Van Inwagen, P. (1990). *Material Beings*. Ithaca: Cornell University Press.
- Varela, F., Rosch, E., & Thompson, E. (1991). *The Embodied Mind*. Cambridge, MA: The MIT Press.
- Wheeler, M. (2005). *Reconstructing the Cognitive World*. Cambridge, MA: The MIT Press.
- Wilson, R. A., & Clark, A. (2009). How to situate cognition: letting nature take its course. In P. Robbins & M. Aydede (Eds.), *The Cambridge Handbook of Situated Cognition* (pp. 55-77). Cambridge: Cambridge University Press