An Exploration of Representational Complexity Via Coupled Oscillator Systems

Anthony Chemero (achemero@phil.indiana.edu)
Douglas Eck (deck@cs.indiana.edu)
Cognitive Science Program
Indiana University
Bloomington, IN 47405

Abstract

We note some inconsistencies in a view of representation which takes *decoupling* to be of key importance. We explore these inconsistencies using examples of representational vehicles taken from coupled oscillator theory and suggest a new way to reconcile *coupling* with *absence*. Finally, we tie these views to a teleological definition of representation.

Introduction

A recent wave of interest in anti-representationalism (Brooks (1991), Beer & Gallagher (1992), van Gelder & Port (1995), van Gelder (1995), to name just a few) has changed the tenor of recent writings on representation: the debate has changed, in part at least, from being about how to determine the *content* of representations to a debate about *what it is to be a representation in the first place*. There have been two primary types of account of representation in this debate: teleological definitions and definitions based on "decoupling". In this paper, we note some inconsistencies in the "decoupling" view of representation and explore these inconsistencies using examples from coupled oscillator theory. We also suggest a way to reconcile coupling with *absence*. Finally, we suggest that these views fit very naturally with a teleological understanding of representation.

Representation and Decoupling

The family of definitions we will be criticizing, those based on "decoupling," derive from John Haugeland's "Representational Genera" (1991). Haugeland's essay has become the touchstone for discussions of anti-representationalist cognitive science, despite the fact that only two of its twenty-seven pages are devoted to discussing representation-in-general. Haugeland's definition, which is in all essentials identical to that of Brian Cantwell Smith (1996)[1], is nicely summarized by Andy Clark (1997) as follows. A system counts as *representation using* just in case:

1. It must coordinate its behaviors with environmental features that are not always "reliably present to the system."
2. It copes with such cases by having something else "stand in" for those features and guide behavior.
3. The "something else" is part of a more general representational scheme that allows the standing in to occur systematically and allows for a variety of related states (Clark 1997, p.144; see also Haugeland, 1991 pp.62-3).

A *representation*, according to this definition, is something that acts as a stand-in in such a system. Without modification, this definition is hopelessly fuzzy. The phrases "stand in" and "reliably present to the system" are left unanalyzed by Haugeland, who only intended this to be "a few dogmatic and sketchy remarks (p.62)."[2] van Gelder (1995), who cites Haugeland's remarks as a source, says that any reasonable characterization of representation will be "based around a core idea of some state of a system, which in virtue of some general representational scheme, stands in for some further state of affairs, thereby enabling the system to behave appropriately with respect to that state of affairs (p.351)." van Gelder, thus, explicitly spells out Haugeland's "standing in" in terms of *behavior*, some state of a system stands in for something else when it enables the system to behave appropriately with respect to that something else. This leaves Haugeland's "not always 'reliably present to the system"' unmentioned.[3] In the remainder of this essay, we will argue that there is no one way to spell out "not always reliably present to the system." In fact, it seems that the phrase is ambiguous among three possibilities, ruling out only those cases in which the representation in the system is constantly causally coupled with the thing represented. This ambiguity has been the cause of considerable confusion. We would be better off if we do not define representation in terms of reliable presence, or lack thereof.

In order to make this case we will need to introduce the following three terms:

- A representation R and its target T are in *constant causal contact* just in case whenever R is present in

---

[1] Smith actually talks about what he calls 'registration', a theory of representation plus ontology. If we deontologize the theory—taking it to be just a theory of representation—it is exactly the same as Haugeland's.

[2] It is rather ironic that Haugeland's "dogmatic and sketchy remarks" have been taken by so many to amount to a definition of representation. In the 1991 paper, Haugeland spends just over a page on defining representation; the paper is 28 pages long.

[3] It is thus unclear whether or not van Gelder takes the "not always reliably present" clause of Haugeland's definition to be important.

a system, T is causing it.

- A target T is *absent* just in case T has no local causal effects when a representation R of it is present in a system. (Both your grandmother and the number 3 are absent in this sense.)
- A representation R is *appropriate* at time *t* just in case R's presence in a system at *t* is not the result of a malfunction of any kind.

Wheeler (1996), who like van Gelder and Haugeland intends only a rough characterization, glosses representation as follows:

> Consider an agent A, one of A's internal states, S, and some feature F, of A's environment (where F may be may be an aspect of A's surroundings, or of A's own activity). A, S, and F are all picked out by some external observer, O. O correctly identifies S to be a representation of F if, on those occasions when O's best explanation of A's behavior is that A is coordinating its behavior reliably with F,
>
> 1. S is the cause of A's behavior, and
> 2. S has been previously (or is now) a cause of A's behavior on an occasion when F was (is) not actually present to A, but—according to O's best theory—A was (is) still coordinating its behaviors with respect to F. (pp.214-5)

Wheeler thus offers (in his second criterion) a somewhat more specified parallel to the Haugeland/Smith[4] requirement that the thing represented need not be reliably present: the representation must cause behavior coordinated with the represented when that represented is absent—that is, distant in space or time—*at least once*. Wheeler's definition would thus be satisfied by a system that occasionally misrepresents. To misrepresent, a system must produce a representation that something is the case, when that something is not in fact the case, and it must do so in a non-appropriate way, such that it is the result of some malfunction in the system. So misrepresentation is one possible gloss of "not always reliably present to the system." Wheeler, then, provides us with a first "not always reliably present": R and T *are not* in constant causal contact, T *is* absent, R *is not* appropriate.

But in *Being There* (1997), which one might call "The Early Clark," Andy Clark claims that the Haugeland/Smith definition requires what he (Clark) calls *decouplability*, which he defines as "the capacity to use the inner states to guide behavior in the absence of the environmental feature [represented](1997, p.144)." This, on the face of it, seems to require no more than Wheeler's definition; that is, the decouplability requirement seems as if it would be satisfied by a system that sometimes misrepresents. Clark apparently does not see it this way, though. He seems to read Haugeland's "not always reliably present" as requiring high-level reasoning: "complex imaginings, off-line reflection, and counterfactual reasoning (Clark 1997, p.147)." Decouplability in the Haugeland/Smith sense, then, is the second "not always reliably present." It requires that R and T *are not* in constant causal contact, that T *is* absent, and that R *is* appropriate. Clark *rejects* decouplability as a necessary condition for something to be a representation. Using a neural group in the posterior parietal cortex of rats as an example, Clark argues that there are systems that are usefully called representational that can never be decoupled from the things about which they carry information. This complex of neurons, he says, carries information about the position of the rat's head, but there is no reason to think that "these neurons can play their role in the absence of a constant stream of proprioceptive signals from the rat's body (p. 145)." So, according to the Early Clark (in *Being There*), something can be a representation when it is always coupled to the thing it represents: a representation, that is, might be in constant causal contact with what it represents.[5]

In more recent papers (Clark, forthcoming; Clark & Grush forthcoming) , however, Clark has made some moves toward reclaiming decouplability. In these papers, a distinction is made between *weak* and *strong* internal representation. A *weak* internal representation is what Clark calls a representation in *Being There*:

> First, [representational mechanisms] must involve inner states whose *functional role* is to coordinate activity with its world (no mere correlations). Second, we must be able to identify *specific* inner states or processes with specific representational roles—we must be able to isolate, within the system, the inner parameters or processes that stand-in for particular extra-neural states of affairs (otherwise we confront only complex inner state implicated in implicated in successful agent-environment coordination) (Clark forthcoming, ms, p.9).

To be a strong internal representation (what Clark and Grush [forthcoming] call an example of *minimally robust* representation), a system or entity must meet a third criterion, in addition to the two above.

> And lastly, the system must be capable of using these inner states or processes to solve problems off-line, to engage in vicarious explorations of a domain, and so on. It is this last capacity that distinguishes model-using agents from the rest. Strong internal representation is thus of a piece with the capacity to use inner models instead of real-world action and search (Clark forthcoming, ms, p.9).

So strong internal representation does require decouplability. In the forthcoming papers, though, Clark attempts to illustrate what he takes decouplability to be in terms of *emulation*. We can call Clark's emulator phase "the Later Clark."

An emulator is a mechanism within a system that takes information about the current state of the system and gives a prediction of the next state of the

---

[4] Not explicitly, though. He does not refer to Haugeland 1991 or Smith 1996.

[5] Indeed, here Clark gestures toward exactly the type of understanding of representation that we will eventually recommend. See Clark 1997, p.146, as well as (Wheeler & Clark, 1997)

system as output (Clark & Grush forthcoming; Grush 1997). Consider (their example) skilled reaching. Moving an arm and hand toward some object depends upon the brain receiving and responding to a stream of visual and proprioceptive feedback concerning the position and trajectory of the arm and hand. But often, due to the inherent speed limitations of the nervous system, the feedback is required faster than it is available. It is in situations like this that emulation is crucial. An emulator in such a case could take as input the current position of the arm and hand, along with the direction of their movement, and provide a sort of mock feedback as output, predicting the position and trajectory of the arm before the actual feedback arrives (Clark and Grush forthcoming, ms, pp.4-5). Emulation, Clark and Grush claim, is important because it is the minimal case of strong internal representation, representation that is decouplable in the Haugeland/Smith sense:

> In sum, it is our suggestion that a creature uses full-blooded internal representations if and only if it is possible to identify within the system specific states and/or processes whose functional role is to act as *decouplable surrogates* for specifiable, usually extra-neural states of affairs. Motor emulation circuitry, we think, provides a clear, minimal and evolutionarily plausible case in which these conditions may be met. (Clark and Grush forthcoming, ms, p.9)

There is some question, however, whether simple cases of emulation really are cases of decoupling in the Haugeland/Smith sense, requiring that the target be potentially *absent.* There is one sense in which we might say that an emulator controlling skilled reaching is decoupled: in the short time between when it takes its input (the state of the arm, the direction of its motion) and when it gives its output (mock feedback to guide action), the emulator is not receiving input directly from what it is representing; that is, unlike the locus of the actual proprioceptive feedback (which is a representation only in the Later Clark's [and Clark and Grush's] weak sense), the emulator's hookup with the target is not constant. This however is *not* decoupling in Clark's original sense: the capacity to use the inner states to guide behavior *in the absence of* the environmental feature represented. Despite the fact that the emulator is not hooked-up to incoming proprioceptive signals for a few milliseconds, the arm and the action it is undertaking is in no way *absent.* With some emulators, that is, R and T *are not* in constant causal contact, T *is not* absent, and R *is* appropriate. But Haugeland/Smith decouplability requires that R and T *are not* in constant causal contact, that T *is* absent, and that R *is* appropriate. Emulation, then, is the third "not always reliably present": R and T *are not* in constant causal contact, T *is not* absent, R *is* appropriate. So emulation is not sufficient for decouplability; thus emulation is not sufficient for Haugeland's requirement that the thing represent need not be always reliably present (at least on Clark's strong reading of it—as requiring "complex imaginings, off-line reflection, and counter-factual reasoning"). So, as depicted in figure 1, claiming that it is a necessary condition on something being a representation that what it represents is "not always reliably present to the system" is ambiguous among three possibilities:

A. It might mean that the system sometimes misrepresents (which would make a comparatively large number of systems representation users);

B. It might mean that the system contains emulators (yielding fewer representations users);

C. It might mean that the system is capable of abstract and/or counter-factual thought (allowing fewer still representation users).

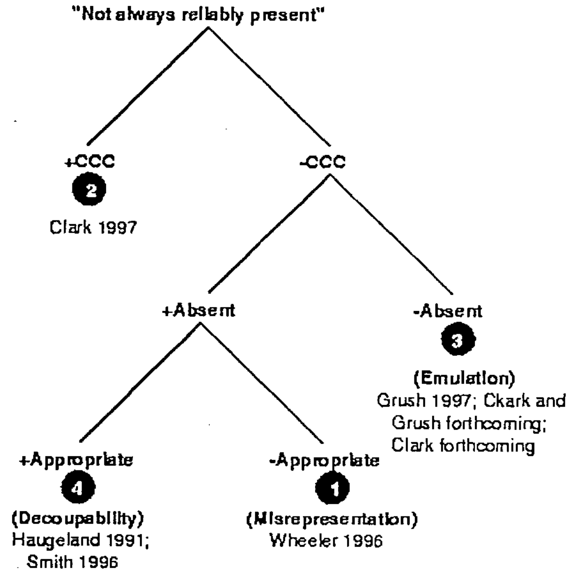

Figure 1: A taxonomy of representational complexity. The numbers are used to refer to oscillator models discussed in the next section.

By breaking down the "now always reliably present" case into substantially different cases we have shown that, without clarification, it is not a workable criterion for whether something is a representation. We now move on to an examination of the refinements offered in the taxonomy (the numbered leaves of the tree in figure 1). We will use mechanisms taken from coupled oscillator theory to uncover some more problems with current *decoupling* views. Specifically, we will show (in order by the numbers on the tree):

1. -CCC, +Absent, -Appropriate (Misrepresentation) is not, by itself, a useful category for understanding representation. Although we agree that every representational system will occasionally represent (a point made by Millikan (1993), we take it that this "ability" is not alone sufficient for viewing a system as representational.
2. +CCC is a useful basic case of representation. Although systems in this category cannot hold onto a representation in the absence of a stimulus, they can be used by a system as an aid in coordinating its behavior with phenomena in the world.

3. -CCC -Absent (Emulation) is not substantially different than +CCC. An emulator is simply a mechanism in constant causal contact with the stimulus, but with a learned time delay.
4. -CCC +Absent +Appropriate (Decoupability) is radically different from the other three, requiring in our example a complex mechanism which is neither neurally or biologically plausible. This is a workable category for some representations but is too restrictive as a necessary condition.

Coupled oscillators as representational vehicles

Coupled oscillators have been suggested as representational mechanisms for a variety of cognitive tasks. McAuley (1995) and Large & Kolen (1994) offer theories of rhythm perception which use oscillators to represent the the relative timing of events. Jones & Boltz (1989) forward a theory of attention which uses an oscillator to represent the level of temporal structure in a task. Even non-temporal tasks like visual feature binding have been modeled using oscillators as the underlying representational vehicle (Singer & Gray, 1995).

A quick review of oscillator theory reveals that two broad classes of biologically-inspired oscillators are often used in cognitive and brain modeling. The first class of oscillators is inspired by *electrical and neural systems*. Several accurate models of neuron action potential (Fitzhugh, 1961; Nagumo, Arimoto, & Yoshizawa, 1962; Morris & Lecar, 1981) come in the form of *relaxation oscillators*, so named because they slowly accrue voltage and then suddenly fire, relaxing or releasing their energy. These oscillators synchronize readily with themselves and with voltage input (Somers & Kopell, 1995). However, a problem for these models as good representational vehicles is that they cannot alone keep hold of a representation in the absence of the causal stimulus. That is, they cannot be +Absent.

The second class of oscillators is inspired by *physical systems* such as mass springs. These models do not synchronize as readily as do relaxation oscillators, partly because they have mass and so have *momentum* keeping them from changing their trajectories to match that of some stimulus. Mass spring systems have been deployed to model many cognitive tasks, though generally those tasks have a motor control component. For example, Thelen & Smith (1994) use mass springs to model the development kicking, stepping and reaching in infants. Also, Schoner & Kelso (1988) use oscillators to model the motor control task of finger wagging as studied by Haken, Kelso & Bunz (1985).[6]

There are of course hybrid models which fall into neither category neatly. These systems tend to take desirable properties from relaxation (electrical) systems and combine them with desirable properties from physical (mass-y) systems. For example the models of rhythm perception by McAuley (1995) and Large & Kolen (1994) use *adaptive oscillators* which synchronize quickly with input signals as do relaxation oscillators but which have something akin to inertia allowing them to "keep the beat" even in the absence of the signal. Gasser, Eck & Port (1996) use a neural network of similar oscillators to learn metrical preferences. These hybrid adaptive oscillators are less stable than relaxation oscillators when coupled together in large groups (Eck, Gasser, & Port, 1999) and are neither physically nor neurally plausible.[7]

Our four coupled oscillator examples are either *relaxation oscillators*, *physical (mass-y)* oscillators or *hybrids*. The relaxation oscillators will handle the cases of +CCC and of -CCC, -Absent (Emulation), categories where Absence is not required. The physical oscillators will handle the "dumb" +Absent -Appropriate case where Absence is required but where Appropriateness is not. We will need a powerful hybrid oscillator—an adaptive oscillator—to handle the +Appropriate (Decoupability) case.

1 – Mass Spring or Pendulum Oscillation (-CCC, +Absent, -Appropriate) Systems which are +Absent and -Appropriate constitute a notably weak class of representational vehicles. These systems lack the necessary adaptation required to avoid "drifting away" from the true state of what is represented.

A simple damped pendulum forced by the input can be seen as such a system. The forcing energy on such a pendulum would be restricted to occur only when the pendulum has a certain phase angle. For example, consider pushing a child on a particular kind of pendulum, a swing set. Presuming the child does not help by changing her moment of inertia at strategic times, the forcing energy will come only from the person doing the pushing and only when the swing is at a particular point in its phase. Otherwise it will be moving wholly in absence of any "stimulus." Furthermore, when the stimulus goes away completely, the pendulum continues moving, allowing it to be +Absent. If the pendulum were equipped with the ability to tap each time it passed some phase zero point we could measure its ability to represent stimuli having temporal organization such as rhythmical patterns. What we would see is that a pendulum can slowly synchronize and beat along with a very restricted set of patterns (those closely matching the resonant frequency of the pendulum) and that it would continue to beat along in the absence of stimulus. However, once the stimulus is removed, the pendulum will immediately begin to drift away from the frequencies found in the pattern back to the exact resonant frequency of the pendulum. In this way the pendulum can be seen as *trying* to maintain connection with the +Absent stimulus, but failing to do so. The stimulus,

---

[6] We're glossing over a small detail here. While Thelen & Smith use true mass spring oscillators in their models, Schoner & Kelso explore the dynamics of moving masses (fingers) with different (non-mass-spring) oscillators.

[7] Note that we are not critical of the adaptive oscillator models of rhythm. They make good predictions about where downbeats occur in rhythmical signals and are considerably simpler than some other candidate models. Given that they were not designed to be used in large groups, their instability in such a context should be forgiven. Furthermore, to our knowledge neither McAuley nor Large & Kolen claim that their model is neurally or biologically plausibility.

that is, is +Absent, but the oscillator misrepresents its frequency.

The forced damped pendulum is described using a simple second-order differential equation found in equation 1. The variable θ is the phase angle of the pendulum. The variable ζ is a damping term. The input, I, is multiplied by $cos(\theta)$ which removes the oscillator from constant causal contact.

$$\ddot{\theta} + \zeta\dot{\theta}|\dot{\theta}| + 4\pi^2 sin(\theta) = cos(\theta) * I \tag{1}$$

The simulation can be made to "tap"by sensing when the pendulum passes its phase-zero point and emitting a pulse:

$$\text{if } cos(x) = 0 \text{ then tap!} \tag{2}$$

2 – Relaxation Oscillation (+CCC) Our second category deals with systems which require constant causal contact to represent. One example of such a vehicle is a simulated neuron such as a Fitzhugh-Nagumo neuron(Fitzhugh, 1961; Nagumo et al., 1962). This oscillator is a simplification of a complete model of neuron action-potential developed by Hodgkin & Huxley (1952) and as we said above it falls into the category of *relaxation oscillator*. When presented with an input pattern consisting of voltage pulses, a Fitzhugh-Nagumo oscillator will synchronize its firing with the pulses. If these pulses are rhythmical, the oscillator synchronizes and "beats along" by virtue of emitting its own pulses in tandem. A connected group of these oscillators can couple with rhythmical input patterns in ways which mirror the *metrical structure* of the patterns. That is, a network can distinguish weak beats from strong beats (see figure 2) and can even represent *rests* using appropriate inhibitory connections.

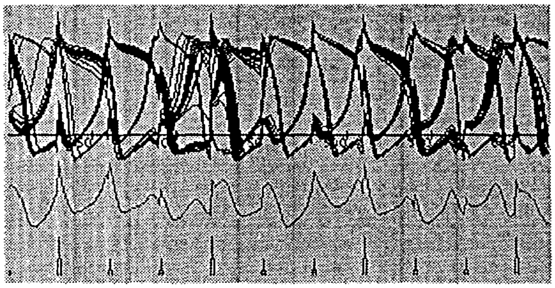

Figure 2: A network of 40 fully-connected Fitzhugh-Nagumo neurons synchronizing with a rhythmical input pattern. The input is at the bottom. Note that despite some instability there are three stable groups of oscillators tracking the three elements of the input pattern consisting of two weak beats followed by a strong beat. The line in the middle is the sum of the fast variable (v) for all the oscillators.

The mathematics of the Fitzhugh-Nagumo artificial neuron is described in equations 3 through 5. The variable v is voltage (or membrane potential) and changes very quickly when the neuron fires. The variable w controls the rate of voltage accrual (or recovery) and changes slowly. The two variables are joined by a coupling parameter ϵ (equation 4) which controls how often the oscillator fires (its period). The oscillator emits a pulse when it fires as measured by the firing point *fp*. The shunt parameter γ is not important for this discussion. Recall that I indicates the input signal.

$$\begin{aligned} \dot{v} &= -v(v-\theta)(v-1) - w + I && (3)\\ \dot{w} &= \epsilon(v - \gamma w) && (4)\\ \text{if } v &= fp \text{ then tap!} && (5) \end{aligned}$$

When the driving stimulus is removed from a network like the one in figure 2 the the oscillators decouple immediately and return to a quiescent state having no relation to the regularities in the input. In this way the oscillators are unable to be Absent. In fact, they simply respond to whatever energy comes to them via constant causal contact. That is, they respond at time t only to the input presented to them at time t.

Yet despite this inability to maintain an appropriate relation to an Absent beat, it seems that simple relaxation oscillators can play the role of a representation in a system. If appropriately connected up within a cognitive system, they can be used by the system to guide its behavior. Indeed, relaxation oscillators play an important role in attempts by one of us to use simulated oscillators control a robotic arm that taps along with beats (see Eck et al., (1999)).

3 – Relaxation Oscillation with Time Delay (-CCC, -Absent) Clark and Grush (forthcoming) provide as an example of emulation a structure in the midbrain of a cat which mimics some of the dynamic parameters of cat forelimb movement.[8] With the use of the emulator, the cat has access to *current* limb position information even though the proprioceptive signal coming from the limbs reports *past* limb position information.

Clark and Grush believe that an emulator like this is special because of its ability to predict current position from past information. We claim that such an emulator is easily implemented using the same representational vehicle as was used above to handle the +CCC case with one minor change. To make an emulating device from a network of Fitzhugh-Nagumo oscillators one need only implement *time delays* so that the network can use the "stale" information from the proprioceptive stream appropriately. Given that cat locomotion is periodic, such a delay could be found simply by measuring the rate of change in the proprioceptive stream and using it to estimate how fast the legs are moving. Once the delay is found, the emulator is able to predict *current* proprioceptive information using information from exactly one leg cycle before. The addition of time delays on the "axons" of the artificial neurons is simple to implement and neurally plausible. (For more on time delays in the propagation of action potentials, see Huber (1998)).

If an emulator is so easy to implement from a relatively dumb network of simulated neurons, why does it seem so *special*? Perhaps it seems special because *in*

---

[8] This example of emulation comes originally from Ghez and Vicaro (1978). See also Kawato, Furukawa & Suzuki (1987).

general it is difficult to predict current states from old information. In difficult tasks like chess playing this is the same thing as predicting what happens a few moves in the future, a task to which considerable computational resources are devoted. If the task did not repeat itself (like cat locomotion), or did not have predictable dynamics (like arm reaching (Thelen & Smith, 1994)) then building such an emulator would be very difficult. But to our knowledge no one suggests that emulators could be used for tasks like this. Even in the domain of motor control no one talks about a *soccer playing* emulator. There is presumably no special emulator in a soccer player's brain which solves the delayed-input problem using *using soccer strategy.* Such a device would indeed be impressive as it would require much more than simple time-delay dynamics, and if emulators existed for such unpredictable or non-repetitive tasks then they would indeed be special. As it stands the emulators suggested by Clark & Grush deal exclusively with repetitive tasks and, we claim, are not interestingly different from relaxations oscillators without time delays. Thus despite claims by Clark and Grush, these are not special representational vehicles after all.

4 – Adaptive Oscillation (-CCC, +Absent, +Appropriate) We present a coupled *adaptive oscillator* as a representational vehicle which handles the case which the Haugeland/Smith view calls *decoupled.* These hybrid oscillators beat along in real time to rhythmical stimuli, a task akin to tapping one's foot along with music. They succeed at this task (called *beat induction*) by taking desirable properties from both physical (mass-y) oscillators and relaxation (electrical) oscillators. See figure 3 for a simulation of the McAuley (1995) adaptive oscillator tracking a slightly noisy metronome.

$$
\begin{aligned}
\ddot{\theta} + (2\pi\Omega)^2\theta &= 0 & (6)\\
I_{windowed} &= g(\theta, I, \chi) & (7)\\
\text{if } I_{windowed} * cos(\theta) &\geq \rho \text{ then} & (8)\\
\theta &= 0 & (9)\\
\text{and } \Omega &= f(\Omega, \Omega_{default}, \theta) & (10)\\
\text{and } \chi &= h(\chi, \theta) & (11)
\end{aligned}
$$

Equations 6 through 11 describe the adaptive oscillator.[9] We will deal separately with the three most important aspects of the hybridization which makes the adaptive oscillator successful: *phase adaptation, confidence rating*, and *frequency adaptation.*

Phase adaptation The adaptive oscillator can instantaneously adapt its phase to match the energy of an input signal (equation 8). When a pulse comes in which is "loud enough" (greater than some threshold ρ) the oscillator immediately resets its phase to zero, as seen in figure 3. In this way, an adaptive oscillator is like a relaxation oscillator which can instantly respond to stimuli but unlike a pendulum or other massy-oscillator, where such infinitely-fast phase resetting is prohibited by momentum.

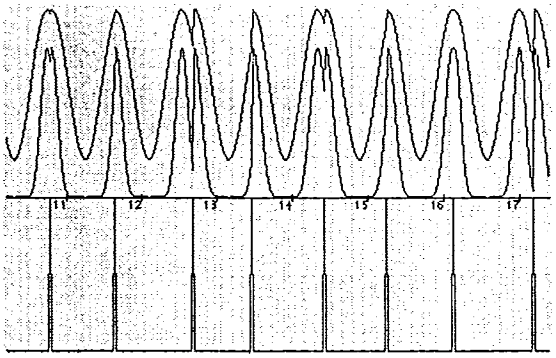

Figure 3: A simulation of McAuley's (1995) adaptive oscillator. At the top of the figure is the phase of the oscillator. In the middle of the figure is a windowing function through which the oscillator filters input. At the bottom of the figure is the input, which is a noisy metronome with a frequency near that of the oscillator's resting frequency. The oscillator can be seen *phase resetting* in an attempt to track the metronome effectively.

Frequency adaptation Performed music and spoken language (two input stimuli which the adaptive oscillator was designed to track) accelerate and decelerate constantly. This poses a problem for a simple phase-adapting oscillator: if the resonant frequency of the input is not matched exactly to the oscillator's resonant frequency Ω, the oscillator will constantly phase reset. This problem is solved by endowing the oscillator with the ability to *tune* its resonant frequency to match that of the input. This tuning is achieved by slightly speeding up or slowing down the oscillator based upon where in its phase cycle it is when a phase reset occurs. In other words, the adaptive oscillator acts like a mass-y oscillator with strong preferred frequency and very much unlike a Fitzhugh-Nagumo relaxation oscillator, which has no voltage-independent preferred frequency whatsoever. We can now see the hybrid nature of the adaptive oscillator: phase resetting is achieved via a relaxation-oscillator type mechanism, frequency adaptation via a tunable mass-spring mechanism.

In equation 10 the function $f(\Omega, \Omega_{default}, \theta)$ speeds up the oscillator if it is consistently phase resetting *after* a beat and slows it down if it is consistently phase resetting *before* a beat. To add stability, limits are placed on how far the frequency Ω can stray from the default frequency, $\Omega_{default}$.

Confidence Rating Without some mechanism other than frequency and phase coupling the oscillator is in constant causal contact with the input (+CCC). That is, at all times the oscillator is sensitive to (loud enough) input. This is a problem for the adaptive oscillator since beat tracking requires that most of the loud sounds be ignored to avoid constant phase resetting. The intuition is that once the oscillator has "found the beat"

---

[9] Though McAuley (1995) and Large & Kolen (1994) offer two adaptive oscillators which differ in the way they modify their instantaneous phase and frequency. For our purposes either oscillator suits hence we make no attempt to tease apart the distinctions here.

it should be relatively difficult to make it phase reset. This problem is solved by adding a phase-dependent window centered at phase zero (equation 7). As the oscillator gets more and more *confident* that it has locked onto the downbeat, it tightens this window (much like putting on blinders so as not to be distracted by events in the periphery) and filters the input $I_{windowed}$. Loudness is no longer enough to cause a phase reset. Now the loud signal must happen at the center of the oscillators attentional window. This phase-dependent window makes the adaptive oscillator a -CCC, +Absent, +Appropriate representational vehicle.

In this way, the adaptive oscillator achieves success by virtue of being +Absent and +Appropriate. More specifically, it is *able* to be appropriate by virtue of being absent sometimes. That is, as it achieves success it ignores more and more of the information in the world and relies more on its own internal state.

Two views of representational complexity For Smith and Haugeland there are two possibilities: a system is either decoupled and representational or coupled and not representational. We believe that it is an error to place these models in two discrete categories. As an alternative we offer a spectrum which moves from *Dumb Coupling* to *Fancy Coupling* as representational complexity increases. See figure 4.

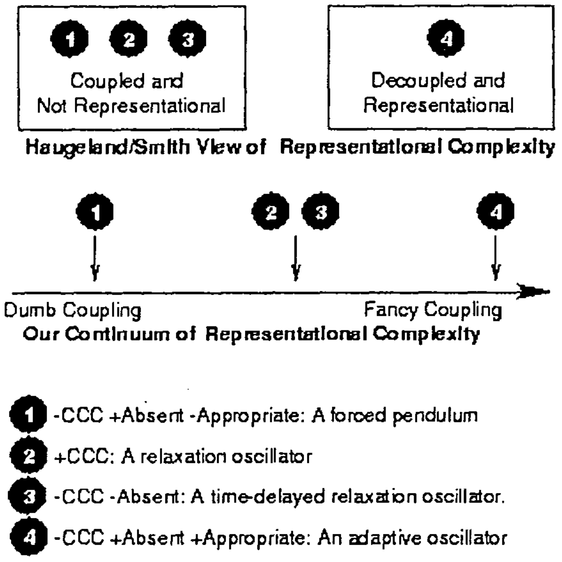

Figure 4: Two ways to categorize representational vehicles in terms of complexity.

The point to take from the Smith/Haugeland view is that a representation must be maintained in some way when its referent in the world is gone. We agree. But we do not believe that this maintenance is best describe by calling it *decoupling.* Our spectrum of *Dumb Coupling* to *Fancy Coupling* gives us a meaningful place for all of these oscillator models. The simple mass spring (Constant Causal Contact) does simple continuous coupling. The emulator is the same: simple continuous coupling with a phase lag. The pendulum which handles the +Absent -Appropriate case does something a little more complex. It couples in such a way that it can ignore the stimulus sometimes. We could call it pulse coupling because the pendulum is only effected by the input when it is at a particular phase location, a coupling regime which looks pulse-like when plotted in time. Finally the adaptive oscillator is an example of fancy coupling. It not only pulse couples but it does so with an adaptively resized window, with infinitely-fast phase resetting and with a constantly- changing adaptive frequency.

A teleological definition of representation

Using coupled oscillators as examples of representational vehicles, we have outlined a spectrum of representational complexity which is based on the "fanciness" of coupling. Such a definition is in conflict with a definition of *decoupling* as held by Smith/Haugeland. However, such a definition is not at odds with a teleological definition, spelled out as follows: A feature R_0 of a system S will be counted as a *Representation for* S if and only if:

- (R1) R_0 stands between a representation producer P and a representation consumer C that have been standardized to fit one another.
- (R2) R_0 has as its proper function to adapt the representation consumer C to some aspect A_0 of the environment, in particular by leading S to behave appropriately with respect to A_0.
- (R3) There are (in addition to R_0) transformations of R_0, $R_1 ... R_n$, that have as their function to adapt the representation consumer C to corresponding transformations of A_0, $A_1 ... A_n$. (See also (Bechtel, 1998), (Millikan, 1984), (Millikan, 1993).)

According to the teleological definition, what some aspect of a system represents, indeed that it is a representation at all, has *nothing to do with what it is or is not coupled with.* Instead, an aspect of a system represents whatever state of affairs (or process) that the representation adapts the representation consuming part of the system to; that is, rather than representing what the world actually is (or isn't) like when the representation is tokened, the representation represents what the world would have to be like for the representation consumer to initiate appropriate behavior. Though, there is not space time to defend it fully here (see Chemero 1998,Chapter 1 for a defense), we will briefly point out a few of its features that are relevant for our purposes here.

Most importantly for our purposes is that this definition of representation allows us to view *all* of these oscillators, with their respective type of coupling, as potential representational vehicles. Unlike the decoupling understandings of representation which draw what we see as arbitrary lines, our definition allows mass spring or pendulum oscillators, relaxation oscillators with or without time delays, and adaptive oscillators to serve as representational vehicles *so long as they meet the three conditions of the definition.* Requiring them to meet these conditions ascribes representations only to somewhat complex systems (with designed mechanisms for producing and using the oscillators as representations); it requires that the representations themselves have particular behaviorally-relevant functions within the system; and it requires that the representations form a i system.

This makes our teleological definition of representation a moderately liberal one—more restrictive than Wheeler's which requires only misrepresentation,[10] but less so than Grush's emulation theory or Haugeland/Smith decoupling. This moderate liberalism is especially important for analyzing arguments in favor of anti-representationalism, the impetus for recent considerations of what it is to be a representation. Anti-representationalism is intended as a radical idea, and indeed it is radical. It should be hard to be an anti-representationalist because, as Rorty (1979) points out, representationalism is a very deeply ingrained part of our entire culture, and not just our attempts to understand the mind. Throwing off centuries of tradition should be no easy matter. Yet if 'representation' is defined in terms of "reliable presence," and hence in terms of "decoupling," anti-representationalism seems fairly timid: even most connectionist networks will not be representational. Defining 'representation' teleologically, as we have suggested, makes anti-representationalism more difficult to maintain and maintains its radical character. To count as supporting anti-representationalism by these lights, a theory or model must be much different than those of mainstream cognitive science, which by now includes connectionism. And that, we think, is as it should be.

References

Bechtel, W. (1998). Representations and cognitive explanations: assessing the dynamicist's challenge in cognitive science. *Cognitive Science*, *22*, 295–318.

Beer, R. & Gallagher, J. (1992). Evolving dynamic neural networks for adaptive behavior. *Adaptive Behavior, 1*, 91–122.

Brooks, R. (1991). Intelligence without representation. *Artificial Intelligence*, *47*, 139–159.

Chemero, A. (1998). *How to be an anti-representationalist.* Ph.D. thesis, Indiana University, Bloomington, IN.

Clark, A. The dynamical challenge. *Cognitive Science*, *Forthcoming.*

Clark, A. & Grush, R. Toward a cognitive robotics..

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

Eck, D., Gasser, M., & Port, R. (1999). Dynamics and embodiment in beat induction. In Desain, P. & Windsor, L. (Eds.), *Rhythm Perception and Production.* Swets and Zeitlinger.

Fitzhugh, R. (1961). Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, *1*, 455–466.

Gasser, M., Eck, D., & Port, R. (1996). Meter as mechanism a neural network that learns metrical patterns. Tech. rep. 180, Indiana University Cognitive Science Program.

Ghez, C. & Vicaro, D. (1978). Discharge of red nucleus neurons during voluntary muscle contraction: activity patterns and correlations with isometric force. *Journal of Physiology*, *74*, 283–285.

Grush, R. (1997). The architecture of representation. *Philosophical Psychology*, *10*, 5–24.

Haken, H., Kelso, J., & Bunz, H. (1985). A theoretical model of phase transitions in human hand movements. *Biological Cybernetics*, *51*, 347–356.

Haugeland, J. (1991). Representational genera. In W. Ramsey, S. S. & Rumelhart, D. (Eds.), *Philosophy and Connectionist Theory.*

Hodgkin, A. & Huxley, A. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*, *117*, 500–544.

Huber, D. E. (1998). The development of synchrony between oscillating neurons. In *Proceedings of the 20th Annual Conference of the Cognitive Science Society.*

Jones, M. R. & Boltz, M. (1989). Dynamic attending and responses to time. *Psychological Review*, *96*, 459–491.

Kawato, M., Furukawa, K., & Suzuki, R. (1987). A hierarchical neural network model for the control and learning of voluntary movement. *Biological Cybernetics*, *57*, 169–185.

Large, E. W. & Kolen, J. F. (1994). Resonance and the perception of musical meter. *Connection Science*, *6*, 177–208.

McAuley, J. D. (1995). *On the Perception of Time as Phase: Toward an Adaptive-Oscillator Model of Rhythm.* Ph.D. thesis, Indiana University, Bloomington, IN.

Millikan, R.· (1984). *Language, Thought and Other Biological Categories.* MIT Press.

Millikan, R. (1993). *White Queen Psychology and Other Essays for Alice.* MIT Press.

Morris, C. & Lecar, H. (1981). Voltage oscillations in the barnacle giant muscle fiber. *Biophysical Journal*, *35*, 193–213.

Nagumo, J., Arimoto, S., & Yoshizawa, S. (1962). An active pulse transmission line simulating nerve axon. *Proceeding IRE*, *50*, 2061–2070.

---

[10] Notice that although the teleological definition has an explanation of misrepresentation built in to it. Since the content of a representation is determined by its function, along with those of the representation producer and consumer, its content will remain constant even in cases in which one or more of the producer, consumer and representation itself fails to work properly. That is, just as the function of a sperm is to fertilize an egg, despite the fact that the number that do so is vanishingly small, so the function of an icon that represents "downbeat-here-now" does not change, even in cases in which the icon is produced or used improperly.

Rorty, R. (1979). *Philosophy and the Mirror of Nature.* Princeton U. Press.

Schoner, G. & Kelso, J. A. S. (1988). Dynamic pattern generation in behavioral and neural systems. *Science*, *239*, 1513–1539.

Singer, W. & Gray, C. (1995). Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience*, *18*, 555–586.

Smith, B. (1996). *On the Origin of Objects.* MIT Press.

Somers, D. & Kopell, N. (1995). Waves and synchrony in networks of oscillators of relaxation and nonrelaxation type. *Physica D*, *89*, 169–183.

Thelen, E. & Smith, L. (1994). *A Dynamic Systems Approach to the Development of Cognition and Action.* MIT Press.

van Gelder, T. & Port, R. F. (1995). It's about time: an overview of the dynamical approach to cognition. In Port, R. F. & van Gelder, T. (Eds.), *Mind as Motion*, pp. 1–44. The MIT Press, Cambridge, Massachuseetts.

van Gelder, T. (1995). What might cognition be if not computation?. *Journal of Philosophy*, *91*, 345–381.

Wheeler, M. & Clark, A. (1997). Genic representation. Tech. rep. 97-002, Washington University, St. Louis.

Wheeler, M. (1996). From robots to rothko. In Boden, M. (Ed.), *The Philosophy of Artificial Life.*