

# Static-Dynamic Hybridity in Dynamical Models of Cognition

Naftali Weinberger and Colin Allen

January 6, 2021

## Abstract

Dynamical models of cognition have played a central role in recent cognitive science. In this paper, we consider a common strategy by which dynamical models describe their target systems neither as purely static or purely dynamic, but rather using a *hybrid* approach. This hybridity reveals why dynamical models should not be understood as providing unstructured descriptions of a system’s dynamics, and is important for understanding the relationship between dynamical and non-dynamical representations of a system.

**Keywords:** dynamical systems, cognitive science, computation, models

## 1 Introduction

Timothy van Gelder’s seminal paper, “What might cognition be if not computation?” (van Gelder, 1995) was an important salvo in the debate between those who take the mind to be a digital computer and those seeking alternative characterizations. As an alternative to the standard computational picture, van Gelder argued that a dynamical systems approach could account for various aspects of real-time cognitive performance, while avoiding various complications and commitments of the computational picture (see also Thelen and Smith (1996) and contributions to Port and van Gelder (1995)). Dynamical modeling of cognitive processes has subsequently become a significant research area and has inspired philosophical developments both in accounts of cognition (Clark, 1998) and of explanation (Bechtel, 1998; Zednik, 2011). In this paper, we argue that these discussions have neglected a key aspect of dynamical modeling that is important for assessing dynamicist claims about cognition.

The conception of computation to which van Gelder was reacting involves machines going through sequences of discrete states in response to discrete inputs. This conception is at the foundation of modern computer science and cognitive science. It was developed into a philosophy of mind via Putnam’s machine-state functionalism (Putnam, 1960, 1967), and more generically as

the computational theory of mind, further elaborated as the computational-representational theory of mind (Fodor, 1981). In developing a dynamical alternative to standard computational accounts, philosophers have emphasized the application of coupled differential equations to model continuous changes through time. Such models reveal how complex patterns of self-regulation can arise from continuous feedback loops in which it is difficult to isolate particular parts as making distinct contributions to the overall behavior. This is claimed to be at odds with the picture provided by standard computational accounts.

Van Gelder’s paradigm example of a dynamical system—the Watt governor—is not itself taken to be a cognitive system, but rather is put forward as providing insights into how cognitive tasks could be performed non-computationally. But which features of dynamical representations make them a promising basis for modeling cognition? It might seem like answering such a question would require an account of what makes a system cognitive, but one can make headway without such an account. It enough to compare the simpler models thus far emphasized by philosophers with dynamical models for the performance of tasks that more closely resemble cognitive ones, in order to see whether the features that have been viewed as significant for cognition are also significant for understanding the performance of the more complex tasks. Towards this aim we will discuss papers by Randall Beer and collaborators (Phattanasri et al., 2007; Beer and Williams, 2015), in which they evolve “minimally cognitive” agents able to perform tasks related to learning and categorization, and then model them using both dynamical and non-dynamical representations. The dynamical models for these agents serve as a basis for evaluating whether philosophical discussions of simple dynamical systems “scale up” to more complex and cognitively interesting ones.

Here we argue that features that have been emphasized by philosophers—specifically those by which dynamical models describe continuous changes and reciprocal agent-environment interactions—are less significant for understanding cognition than other features that have been neglected. In dynamically modeling a system, the longer-term stability properties of the system are typically as important as facts about which quantities are changing, and the representation of certain quantities as either constant or as variable itself depends on decisions about how to model the system. In the following, we will particularly emphasize a standard representational device by which a model will represent the same quantity alternately as a time-varying variable in one equation and as a fixed parameter in another. This device enables the modeler to track the dynamical evolution of a system in the context of a broader “attractor” landscape, such that the instantaneous change in the system’s trajectory at each moment depends on the abstract position of the system within the larger landscape (e.g., how far it is from equilibrium). For now we will refer to representations employing this device as *hybrid* representations. In section 3.3 we will characterize this phenomenon more precisely using the concept of a “quasistatic approximation” (Beer and Williams, 2015).

The dynamical model of the Watt governor and the dynamical models developed by Phattanasri et al. (2007) and Beer and Williams (2015) employ qua-

sistatic approximations. We will argue that as the systems considered become more complex, the use of quasistatic approximations becomes more significant for understanding their ability to perform cognitive tasks. Since all three cases employ representational hybridity, it is clear enough that such hybridity alone is insufficient for describing a system as cognitive. Whether a system performs a cognitive task cannot be read off from the formal features of its dynamical models. Nevertheless, we will describe several ways in which our discussion of hybridity leads to a better understanding both of dynamical models and of their relationships to non-dynamical models. To start, it helps dispel a naive picture on which dynamical models provide an spatiotemporally unstructured characterization of a system. This in turn eliminates a worry that any discretization or modularization of a system must be arbitrary, since a proposed way of carving up a system can be justified by its tracking features already present dynamics. Finally, it makes salient an important subtlety in the way that dynamical models represent their target systems. Many of the choices going into developing a standard computational model—such as the division of a system into functional modules—are easy to see in the model itself. But dynamical models are just as much representational devices as standard computational models, even if they do not wear the modeler’s representational choices as much on their sleeves.

## 2 Representations, Mechanisms and Explanations

Van Gelder’s (1995) paper has been very influential among philosophers working on a variety of issues, from metaphysics of mind (e.g., whether minds are representational, computational, etc.) to the nature of explanation in cognitive science (e.g., whether dynamical models fit criteria for being explanations). Van Gelder deploys the Watt governor for regulating steam engines as a metaphor for the causal relationships operating in brains and the bodies that contain them. The Watt governor consists of a spindle with two hinged arms whose rotation is coupled to a steam engine. As the rotational speed of the spindle increases or decreases, the arms go up or down, respectively closing or opening a valve (i.e., a throttle) that controls the flow of steam, thus regulating the speed of the engine. This negative feedback loop serves to stabilize the otherwise rather erratic behavior of steam engines that were due to factors such as fluctuations in the heat produced by burning coal and changes in the external load on the engine. The mathematical analysis of governors took considerable work by Maxwell and others, to develop a set of continuous differential equations whose shared parameters model the couplings across the system.

Van Gelder argues that standard computational theories with their commitment to internal symbolic representations are inadequate to describe the workings of the engine-governor system. He emphasizes the continuous and reciprocal nature of the causal interaction between the angle of the arms and the speed of the engine, and he claims that this relationship is “much more subtle and complex than the notion of representation can handle” (1995, p.353). Instead, he maintains, this framework requires the “mathematical language of

dynamics” (ibid.) within which these quantities are coupled. He concludes this discussion by stating: “The real problem with describing the governor as a representational device, then, is that the relationship of representing—something standing for some other state of affairs—is too simple to capture the actual interaction between the governor and the engine” (ibid.).

It is worth highlighting van Gelder’s inference from the claim that a particular relationship is described using the mathematical language of dynamics, to the claim that any computational-representational model would be too simple to capture the actual interactions. But what is the relationship between this mathematical language and the target system it represents? To say that computational-representational models are too simple to capture the actual interactions in the system is to presuppose that there is some minimal standard for what counts as adequately representing the system. But from the fact that one can represent interesting features of the governor’s behavior using models from dynamical systems theory, it does not follow that one must represent those features in order to adequately represent the system for some purpose. Neither does it follow that the differential equations by themselves provide the minimal standard.<sup>1</sup>

Van Gelder’s use of features of the dynamical representation to argue that the Watt governor cannot be understood on a standard computational-representational model is most salient in the following passage:

[In the Watt governor, n]ot only are there no representations to be manipulated, there are no distinct manipulatings that might count as computational operations. There are no discrete, identifiable steps in which one representation gets transformed into another. Rather, the system’s entire operation is smooth and continuous; there is no possibility of non-arbitrarily dividing its changes over time into distinct manipulatings, and no point in trying to do so. (van Gelder, 1995, p. 354)

In focusing on smoothness and continuity, van Gelder is appealing to mathematical features arising in the application of differential equations. We will argue that the question of whether a dynamical system can be non-arbitrarily decomposed cannot be resolved by appealing to such mathematical features.

Aside from smoothness and continuity, van Gelder also emphasizes that the governor’s activities “are happening continuously and at the very same time” (van Gelder, 1995, p. 354). While this is not an inaccurate description of the system, it neglects important subtleties in the way that time is represented in dynamical models, as we will explain. Grush (1997) has already noted a mismatch between the temporal features highlighted by proponents of the dynamic approach and those of the models they consider. In particular, while proponents emphasize “real-time” dynamics, many models used to demonstrate the

---

<sup>1</sup>In criticizing van Gelder for not attending to the relationship between the dynamical *representation* and the target system, we are not directly addressing his position that the target system should not be understood as processing *representations*. We flag the use of these two notions of representation to prevent possible confusion.

promise of the approach are not in real-time, but involve discrete time-steps corresponding to task-sequences.

Before proceeding, it will be useful to differentiate the questions we will be addressing from those that have been considered elsewhere. There has been considerable discussion in the philosophical literature about whether dynamical cognitive science excludes representational accounts of mind (Bechtel, 1998) and whether it provides an alternative view of computation or is incompatible with computational theories of mind (e.g., Wheeler, 2005). Authors in this debate have been concerned with such questions as whether dynamical models are genuinely explanatory or merely descriptive of target systems, and if explanatory, whether they conform to patterns of mechanistic explanation, causal explanation or something else (see, e.g., Clark, 1998; Bechtel, 1998; Wheeler, 2005; Chemero and Silberstein, 2008; Wilkenfeld, 2014; Kaplan, 2015).

In this paper, we do not engage with the vast debate on the nature of scientific explanation in general or of model-based explanation more specifically (e.g., Bokulich, 2011). Nevertheless, we will briefly illustrate how participants in this debate have largely focused on the features of dynamical models emphasized by van Gelder. For instance, debates over the explanatory status of dynamical models have concerned their purported inability to explain a phenomenon by decomposing its mechanism into localized components (Bechtel, 1998; Chemero and Silberstein, 2008; Kaplan and Craver, 2011). These debates direct one's attention to features of dynamical models that supposedly threaten localization. This is just one way in which van Gelder's emphasis on decomposition has had long-lasting influences.

Zednik (2011) provides a good example of the extent of this influence. Zednik has argued that dynamical explanations in cognitive science, despite sharing the common feature of being formulated via differential equations, do not, in fact, constitute a single explanatory type. Some models (e.g., those of Thelen and Smith (1996) and Beer 2003) do, he maintains, support decomposition into parts and their activities that allows them to be characterized as providing mechanistic explanations, whereas others (e.g., Haken et al., 1985) do not, rather providing a covering law explanation (Hempel and Oppenheim, 1948; Bechtel, 1998).

Zednik is correct to highlight the plurality of dynamical models and to put pressure on hasty arguments for why such models cannot provide mechanistic explanations. But he is uncritical in his assumptions about what it is that models of dynamical systems must explain. He considers two primary challenges for modeling certain dynamical systems mechanistically. One is that dynamical models involve an agent-environment interaction. The other is that coupled dynamical systems involve continuous reciprocal causation and thus allegedly cannot be decomposed into localized parts with localized functions. Regarding the second claim, Zednik assumes that for systems involving reciprocal continuous causes it will be "difficult or impossible to allocate responsibility for any particular operation to one part of the system" (p.259), and thus must be represented using the models of dynamic systems theory. He merely disputes that this entails that dynamic models cannot be given a mechanistic interpretation.

For all that he moves the discussion forward by distinguishing different types of dynamical models, he nevertheless accepts van Gelder’s characterization of the features of dynamical representations that merit philosophical discussion.

Our criticisms notwithstanding, philosophers since van Gelder have considered a wide range of dynamical models that are significantly more complex than that for the governor. A notable example is Eliasmith (2010), who posits an important cognitive difference between dynamical systems that can and cannot be modeled using the tools of control systems theory. Nevertheless, our discussion in this section motivates a more general discussion of the features of dynamical systems that are relevant to cognition. One advantage of focusing on Beer and colleagues’ “minimally cognitive” agents is that the agents are evolved without making any a priori assumptions about how they should perform the task. They thus provide a good basis for an empirically grounded discussion of which modeling frameworks are suitable for modeling their dynamics.

### 3 Hybridity in Dynamical Representations

Van Gelder and subsequent writers have emphasized the use of differential equations to represent the continuous evolution through time of a system of closely interacting parts. This emphasis is the result of focusing on the derivatives in differential equations, which are well defined for all values of a function when that function is smooth (in the sense of being everywhere differentiable). Yet a myopic focus on derivatives can lead philosophers to miss the range of modeling decisions that go into modeling a system dynamically. These decisions include: the specifications of initial conditions, boundary conditions, and rigidity constraints, as well as the choice to model certain quantities using variables and others using time-invariant parameters. These decisions reveal that dynamic modeling is not merely a matter of specifying what quantities are changing, but also which remain stable over a time-period of interest.

In this paper we focus on the way that dynamical models rely on assumptions about the equilibrium or attractor states of a system, and how such assumptions are employed in modeling the system’s evolutionary dynamics. In the models we will consider, the dynamic models cannot be understood as “pure” representations of a systems dynamics, but rather as “hybrid” representations in assumptions about the longer-term stability of a system play a role in modeling the shorter-term transient dynamics. After highlighting a few examples of such hybridity, we will show how it is more rigorously characterized through the concept of a *quasi-static* representation (Beer and Williams, 2015).

In this section we will provide three examples of hybrid representations. The first is based on a closer inspection of van Gelder’s treatment of the Watt governor. The second and third are from more recent work by Randy Beer and his collaborators. By beginning with van Gelder, we aim to show that the modeling device we are describing is employed (and typically ignored) even in widely-discussed examples from the literature. Crucially, in highlighting formal similarities across the three examples, we are not suggesting that the examples

are all similarly relevant for modeling cognition. The behaviors modeled for the agents in the last two examples are more complex and relevant to cognition than that of the Watt governor. Nevertheless, our focus on hybridity provides a useful lens through which to compare the models. In the section 4 we will further discuss the relevance of hybridity for cognition, and will explain why it becomes even more significant as one considers models for more complex cognitive behaviors.

While here we focus on cognition, it is important to note that the features of dynamic models we identify are ubiquitous across the sciences. The philosophers Jordi Cat Cat (2005), Mark Wilson Wilson (2017), Sarah Green and Robert Batterman Green and Batterman (2017) have been particularly attentive to the way that dynamical models in areas as diverse as physics and biology use subtle representational devices to incorporate information about a system’s steady-state and equilibrium behaviors. Accordingly, a virtue of the present discussion is that it creates a potential bridge between work on dynamical systems in cognitive science and the more general study of how dynamical models function across the sciences.

### 3.1 Modeling the Watt Governor

We have already described the basic structure of the Watt governor above as consisting of a spindle with flywheel arms connected to a throttle controlling the amount of steam flowing in the engine. The dynamical model of the governor represents the behavior of the system using terms that explicitly represent the angle of the arms from the vertical, the speed of the engine, and the throttle setting. The governor’s key dynamical feature is the feedback loop by which it regulates the speed of the engine so that the speed does not substantially deviate from a desired value. The dynamical model captures this through linked equations sharing the common terms just mentioned. Although Maxwell’s model is more complicated, van Gelder, following Beltrami (1987), boils it down to a pair of equations (or, more accurately, one equation and one schema for an equation).

The first of these equations describes the current acceleration of the angle of the arms given the current value and velocity of the angle.

$$(1) \quad \frac{d^2\theta}{dt^2} = \underbrace{(n\omega)^2 \cos(\theta) \sin(\theta)}_{(i)} - \underbrace{\frac{g}{l} \sin(\theta)}_{(ii)} - \underbrace{r \frac{d\theta}{dt}}_{(iii)}$$

Here  $\theta$  is the angle of the arms,  $\omega$  is the speed of the engine and  $n$ ,  $g$ ,  $l$ , and  $r$  are constants. The acceleration of the angle is given as a function of three terms (which we have labeled (i), (ii), and (iii) above), the first two of which involve the current angle of the arms ( $\theta$ ) and the third involves its velocity ( $d\theta/dt$ ). The current value of  $\theta$  influences its acceleration by determining the outward effect of the force exerted by the engine (term (i)) and the inward effect of gravity (term (ii)— $l$  is for arm length). The velocity of  $\theta$  influences the acceleration of

$\theta$  by producing friction at the hinges of the device linking the governor to the throttle valve (term (iii)). This dampening influence of friction is necessary for the system to stabilize (just as without friction or air resistance a pendulum will continue to swing indefinitely). The construction of the model is based on Maxwell’s original work in which he explicitly invoked the notions of kinetic energy, potential energy, and friction (or resistance) corresponding to the three terms respectively, adapting the equations for a pendulum (terms (i) and (ii)) that is damped (term (iii)).


As van Gelder notes, in equation (1)  $\omega$  is treated not as a time-dependent variable, but as a fixed parameter. This may seem puzzling, since it is crucial to the functioning of the governor that the speed of the engine changes as a function of  $\theta$ . In discussing equation (1), van Gelder temporarily considers the case in which the governor is detached from the throttle valve so that the engine speed no longer depends on  $\theta$ . Under such a scenario the engine speed could be constant, but we still need an explanation for why equation (1) remains applicable to the case in which the link between the governor and the throttle valve is not broken. Here the key is to focus on the role of equation (1) in predicting the stability of the system. Think of (1) as providing a snapshot of the system at a time. Does the snapshot represent the system at a stable equilibrium point? We can determine this by imagining that the acceleration and the velocity of the angle equal zero—as they would when the system is at steady state. Doing so reveals that the system will be at equilibrium only when the first and second terms are equal. Additionally, whether the difference between these terms at points near equilibrium is positive or negative in the neighborhood around equilibrium will determine whether the equilibrium point is a stable one.

While the engine speed  $\omega$  is given in equation (1) as a parameter (understood in this context to be a non-time-dependent variable), in the second equation it is modeled as a time-dependent variable. This is essential for modeling it dynamically, since if it were modeled as a constant rather than as a variable, then its derivative would be zero at all times. Van Gelder presents the second equation for the influences on the derivatives of  $\omega$  in the dynamical model schematically:

$$(2) \quad \frac{d^n \omega}{dt^n} = f(\omega, \dots, \tau, \dots)$$

where  $\tau$  is the setting of the throttle valve. Van Gelder does not fill in the details of this equation (e.g., the order of the derivative or the additional variables in the function) and we are willing to grant him that for the sake of modeling the feedback loop in this system that these details do not particularly matter – i.e. they can be safely black-boxed. Yet the fact that  $\omega$  is represented alternately as a parameter and as a time-dependent variable is important for understanding the representational division of labor underlying the dynamical model. In modeling  $\omega$  as a parameter, one represents the change in the acceleration of  $\theta$  resulting from the current state of the system, in particular how far the system is out of equilibrium. One can then take the individual snapshots of the instantaneous influence of  $\omega$  on the acceleration of  $\theta$  and combine this with information about





**Figure 1:** Dynamic Causal Model for the Watt Governor: Solid arrows indicate causal arrows, dashed arrows (labeled “I”) indicate integration links (see text for details)

how  $\omega$  varies as a function of  $\theta$  (via the throttle setting) in order to model the evolution of the system. This requires modeling  $\omega$  using a time-varying variable.

A useful framework for modeling self-regulating systems such as the governor is provided by Iwasaki and Simon’s dynamic causal models (Iwasaki and Simon, 1994). These models take a causal modeling framework that was designed for static sets of equations giving the causal relationships between simultaneous variables, generalizing the framework to model a system in which some of the variables are away from equilibrium. Here we won’t dwell too much on the causal interpretation of these models (Rescher and Simon, 1966; Dash and Druzdzel, 2001; Weinberger, 2019), but will use them primarily as a way to keep track of the temporal relationships among the variables for the governor.

In figure 1 we have applied Iwasaki and Simon’s technique to provide a visualization of the relationships among the differential equations in the model of the governor. The influence of  $\theta$ ,  $\theta'$  and  $\omega$  on  $\theta''$  is read off of (1) and the influence of  $\tau$  on  $\omega$  is given by (2). We have also added a causal arrow from  $\theta$  to  $\tau$ , as the influence of the governor on the engine speed is essential to its functioning. Solid causal arrows model “simultaneous” relationships, and dashed arrows, or integration links, correspond to the mathematical operation of integration—that is, of taking the integral of the derivative function. The so-called simultaneous relationships in the model need not be taken as entailing that the causal influences represented in fact take no time. Rather, they may be understood as indicating that the effect variable has had sufficient time to respond to any changes in its cause(s) at the point at which both variables are measured. For instance, the mechanical coupling between spindle arms and throttle is not modeled dynamically, with the underlying assumption being that this connection is effectively rigid enough to be treated as instantaneous given the modeler’s interest in what happens at the given time scale. This contrasts with variables that are linked via derivatives and integration-links. Through integration, one can take the value of a variable at a time-step and give a discrete approximation of its value at the next time-step. The integration links are from higher-order derivatives to lower-order derivatives. All of the solid causal arrows are straightforwardly derived from interpreting equations (1) and (2) as equations in which the variable on the left-hand side is an effect of the variables on the right-hand side.

One useful feature of the representation in figure 1 is that it enables one

to easily check that a necessary condition for the system’s stability is met. Specifically, for a system to be stable, the highest-order derivative of any variable must be a function of the variable’s lower-order derivatives. If, for example, the angular acceleration  $d^2\theta/dt^2$  did not change with velocity, it could not be a feature of the dynamics that it would be pushed back towards zero once the arms were in motion. It is straightforward to see that the dependence of the highest order derivatives on the lower orders is met in the present model (although this alone does not guarantee stability).

In figure 1,  $\theta$ , but not  $\omega$ , is modeled along with its time-derivatives. This corresponds to the same division of labor involved in treating  $\omega$  alternatively as a parameter and as a variable. Representing  $\theta$ ’s velocity and acceleration enables one to capture the feedback loop by which the system as a whole tends towards a constant speed. Including derivatives for a variable enables one to represent that variable (in this case  $\theta$ ) as having been perturbed from a stable state and as not having had adequate time to return to that state. Note that changes to a variable’s velocity or acceleration at a time do not change that variable’s value at that time, although they will influence that variable’s value an arbitrarily short period of time later (as can be calculated using integration). In contrast,  $\omega$  is represented without a time-derivative, and thus *as if* it responds instantaneously to any change in the value of  $\theta$ . As long as  $\theta$  has not reached its long-term steady-state value, neither will  $\omega$ , but the model attributes the system’s being away from equilibrium to the “stickiness” of  $\theta$ . Given enough time for  $\theta$  to adjust to prior perturbations  $\omega$  can be treated as immediately following suit.

The dynamic model for the Watt governor highlights the way that the coupled differential equations for the system do not provide a pure unstructured description of its dynamics, but are specifically designed to model its equilibrium-preserving behavior. This is done via a hybrid form of representation in which  $\omega$  is alternately treated as a parameter and variable in the two equations. We now turn to more complex dynamical models involving similar hybridity.

### 3.2 Phattanasri and Beer

We now turn to two papers by Randall Beer and his collaborators coming out of their sustained effort to understand and defend the application of dynamical models in cognitive science. We will focus on Beer’s projects in which he evolves simulated neural network agents to perform relatively simple tasks related to learning and categorization. These agents’ neural networks are constituted by sets of dynamical models in which the parameters are tuned through a genetic algorithm simulating the natural selection of the agents over many generations, where fitness is defined in terms of their abilities to perform the relevant task. The systems developed barely register as “cognitive”—they are, in Beer’s parlance, “minimally cognitive agents”. Nevertheless, their capacities and their dynamics are considerably more complex than the centrifugal governor. While the governor has been seen as providing insights into cognition, few have suggested that it is in fact a cognitive system.

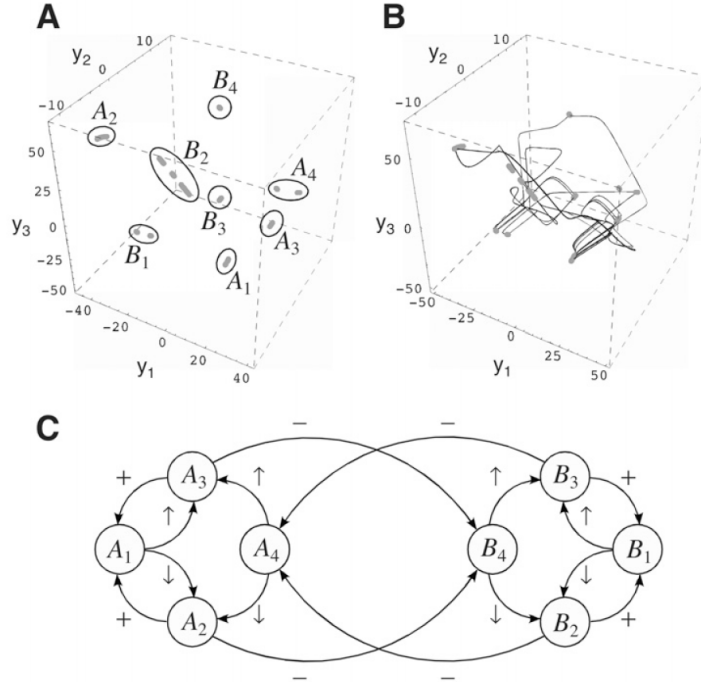
This subsection considers an experiment by Phattanasri et al. (2007) aiming to understand the dynamics of an evolved artificial agent selected for its ability to adjust its behavior to a contingent, changing relationship between a cue stimulus and a reward. In the experiment, the cues were labeled as “smells”, which were predictive of two kinds of “food”, but where the relationship between the smells and the food was reversed unpredictably. Artificial selection was applied to simple agents with up to six internal neurons on the basis of whether they successfully timed the opening of their “mouths” to obtain the “edible” food, or kept their mouths closed to avoid ingesting “inedible” foods. Because the relationship between cues and the positive or negative reinforcement provided by edible and inedible food was reversed during the agents’ “lifetimes”, successful agents following a variable response strategy had to change behavior as a result of experience (a capacity Phattanasri et al. refer to as “learning”).

Phattanasri et al. focus their analysis on the simplest successful agents, which had three neurons. To study the internal dynamics of a successfully evolved agent in this task, they applied a couple of techniques: (A) They tested the agent under conditions that were not part of its “evolutionary history”, such as clamping an input to the network (holding it constant) as if a smell cue appears but does not disappear. This enables them to identify attractor basins in the state space defined over the activation levels of the agents’ three neurons, which they represent using “phase portraits” (not pictured here, see Phattanasri et al., 2007, 386). (B) They then map a trajectory of the non-equilibrium dynamics by which the agent’s position in the state space changes as a result of new inputs such as smells or positive/negative reinforcement (see figure 2).

The two-step process just described involves hybridity analogous to the dual treatment of  $\omega$  in the two equations for the Watt governor. Just as with the governor, where the initial treatment of  $\omega$  as a parameter was not to indicate that  $\omega$  is constant, but rather as a basis for then calculating how the system changes as  $\omega$  changes exogenously, Phattanasri et al. begin with these phase portraits in order to then consider how the neuronal states will evolve as the inputs shift. In both cases, this representational division of labor should not be taken to indicate that certain parts of the system are stable and others varying, but rather to capture both equilibrium and non-equilibrium dynamics within a single representational framework.

Because the exact behavior of the agent varies from trial to trial, Phattanasri et al. use a “strobing” technique to build a composite picture of the non-equilibrium dynamics over multiple trials, showing that the agent’s states tend to cluster in localizable regions of the state space just after key events such as the appearance of a cue or the appearance of a reinforcer (Phattanasri et al., 2007, 388)(see part A of figure 2 below).

Using the results of these techniques, they identified regions of state space in which the systems tend to be found under actual and possible input conditions. By mapping and linking the strobed regions of state space which the system tends to be passing through at critical points during the task, and treating those regions as states of a finite state machine (FSM), they constructed a



**Figure 2:** Figure from Phattanasri et al. p. 388. The axes  $y_1$ ,  $y_2$ , and  $y_3$  refer to the states of neurons 1, 2, and 3, respectively. The system was strobed at the end of each smell and reinforcement signal. Figure (A) uses ovals to divide the strobed regions into distinct clusters. (B) depicts a trajectory through the state space in a trial involving multiple changes of environment. (C) provides a finite state machine representation for how the system responds to signals with  $\uparrow$  and  $\downarrow$  denoting the two different “smells” and  $+$  and  $-$  denoting to positive and negative reinforcement signals.

FSM representation that switches between two different cycles depending on which of two “smell” cues is currently predictive of the “edible” reward. They variously refer to this FSM as “embedded” (Phattanasri et al., 2007, 388, fig. 7) in the evolved neural circuitry and “extracted” (ibid.) from the dynamics, and they go on to explain that the circuits “work by implementing finite state machines that capture the sensation-action-reinforcement structure of this task” (Phattanasri et al., 2007, 391). Below we will provide some critical discussion of the precise relationship between the dynamics and the FSM representations, and will argue that the hybridity of the representation is relevant to understanding this relationship.

### 3.3 Beer and Williams on Quasistatic Approximations

In a paper that has implications both for the present discussion as well as the debate that motivates it, Beer and Williams (2015) compare approaches to cognition using information theory (IT) to those using dynamic systems theory (DST). Beer and Williams analyze a set of artificial agents who were evolved to make a behavioral decision (intercept or avoid) based upon an asynchronous comparison of the relative size of two objects which move towards the agent. While DST approaches track the evolution of the agents' neuronal activity, IT provides tools for determining how information about the sizes of the cue and probe is distributed through the neurons over time. Beer and Williams view the IT and DST representations of the agents not as competitors, but rather complementing one another, and we are sympathetic to this position.

Before presenting the details of the experiment, we will first convey their notion of a *quasi-static approximation*, which provides a more rigorous way to understand the form of representational hybridity we have been discussing. Beer and Williams explain how one and the same dynamical system can be represented either as a single system governed by a time-invariant set of equations, or as two or more coupled systems in which the output of one determines the values of the parameters in the equations of the others at a given time. While in the former case, the dynamical law governing the system is autonomous, meaning that it is fixed, in the latter case, the laws governing each subsystem are non-autonomous, meaning that the parameters in their dynamical laws change over time. The autonomous and non-autonomous perspectives can be combined into a *quasistatic approximation* (Beer and Williams, 2015, 13) in which one represents the non-autonomous dynamics of a part of the system as the result of a series of snapshots, in each of which the dynamics are treated as autonomous. This is what was going on in the Beltrami/van Gelder model of the Watt governor, where  $\omega$  in equation (1) was treated as a fixed input rather than as time-varying, and thus *as if* the dynamics were autonomous. Phattanasri, Chiel and Beer similarly rely on a quasistatic approximation when they appeal to phase portraits describing the dynamics of the system given fixed inputs, prior to using these autonomous representations of the dynamics to account for the system's transitory non-autonomous dynamics when away from attractor states (Phattanasri et al., 2007, 384).

Beer and Williams evolved agents to perform a task consisting of two stages. In the first stage, the agent passively observes a falling "cue". In the second stage, the agent must either catch or avoid a falling "probe" depending on whether it is bigger or smaller than the cue. In their dynamical representations, the state space of the agent maps the relationships among the activation levels of the agent's different neurons as a function of time or a proxy of time. Different trajectories correspond to different trials with different cue sizes. One of their aims in dynamically modeling the systems is to understand the different trajectories of the system in the cases where the agent either catches or avoids the probe.

The bundles of trajectories in which the agent either avoids or catches the

probe correspond to different attractor states in the dynamical landscape. Beer and Williams (2015, 16, figure 7) provide several graphs displaying the different bundles and pinpointing when they diverge towards the different attractors. Yet, as they emphasize, it is crucial not to think about the attractor landscape as a fixed map through which the trajectories travel—the attractor landscape changes both simultaneously with and in response to the changing trajectories of the neurons through the state space. In the early moments of the probe’s descent all of the trajectories are all heading towards a single attractor point, but over time the attractor landscape changes to one in which there are two distinct attractors towards which the different bundles tend. Such a change in the attractor landscape is called a *bifurcation*.

Beer and Williams’ use of bifurcation diagrams to explain the split in the trajectories involves a use of the quasistatic approach (Beer and Williams, 2015, p.13). Both the agent and its environment change as a function of time, and in principle one could model the agent-environment system as a single dynamical system with a fixed dynamical law. Instead, Beer and Williams model the agent and its environment as two dynamical systems with dynamical laws that change over time. While the agents’ neuronal dynamics are constantly changing in response to the changing sensory inputs, at each point in time the sensory input is represented as fixed. In this manner, one can represent the way that the agents’ states change in response to the sensory input without explicitly representing how the sensory inputs change over time. This representational choice makes it possible to visualize a bifurcation diagram in which one can transparently represent both how particular trajectories change over time and also how different trajectories fall into different attractor basins. While the sensory inputs are also changing over time in response to the agents’ movement, the quasistatic approach enables the modeler to model this independently and to thus gain an understanding of the splitting of the trajectories that would be unavailable otherwise.

The way in which Beer and Williams take the time-dependent variable for the sensory input and model it as a parameter in the agent’s differential equations is, from a formal perspective, identical to the way in which  $\omega$  functions alternately as a variable and a parameter in the equations for the Watt governor. Yet the behaviors illuminated by the use of the quasistatic approach are much more interesting in the former case than in the latter. Now that we have illustrated the common use of the approach in the three different cases, we now use this commonality to illuminate the features of Beer’s agents that make them more suitable than the governor as models for cognitive processes.

## 4 Quasistatic Models of Cognitive Processes

The Watt governor is, of course, entirely boring from a cognitive perspective; because of the design of the system it can only be just slightly out of equilibrium, and it is always tending towards a fixed equilibrium point in the absence of further input or inherent minor fluctuations in state. While the “snapshots” of

the governor employed in the quasistatic approximation are informative about how far the system is away from equilibrium, and thus about how long it will take to return to equilibrium, such snapshots will at most provide a cumulative record of the prior perturbations to the system, with no way of distinguishing between different types of perturbations—e.g. a decrease in the total workload as opposed to a change in the combustion driving the engine.

Phattanasri et al.’s dynamical models are more complex in that they describe systems with multiple attractor states. Different regions of the state space in the vicinity of different attractors are associated with different responses to the same smell, and the agent’s position in the state space shifts from one region to another in response to negative reinforcement signals. While the purpose of the governor is to make the behavior of a device relatively invariant to changes in its environment, the task performed by Phattanasri et al.’s agents requires them to change their behavior based on signals they receive from their environments. The position of the agent near a particular attractor in the state space is thus discriminative among different possible causal histories—e.g. whether the agent is in an environment where a particular smell is to be pursued or avoided. It is important to realize, however, that it is not important—or optimal—that that the system ever settles into a particular attractor, and its doing so would in fact hinder the agents’ abilities to respond quickly and adaptively to subsequent stimuli.

As a result of the increased complexity of Phattanasri et al.’s dynamical models, as compared to the Watt governor, the “snapshots” of the quasistatic approximation are more informative. The location of the agent in the abstract state space relative to an attractor state provides information about its causal history and thus of the environment that it is in. Representations of the system at a single point in time are thus both informative about its transitory and longer term dynamics.

The dynamical models for the agents developed by Beer and Williams (2015) are even more complex, in the sense that the dynamical attractor landscape changes over time and involves a bifurcation. Although in all of the cases where quasistatic approximations are employed, it is important to realize that the parameters that are treated as unchanging at a time are not in fact unchanging, this feature is especially crucial for understanding the activities of these agents in catching or avoiding the probe. While it would be possible to model the agent and its environment as a single system in which both are constantly changing as a function of a single dynamical law, the use of the quasistatic approximation enables one to model the agent’s internal dynamics (semi-)independently of the broader attractor landscape in a way that (as we will further explain) enables one to better understand its decision-making behavior.

All three models considered are dynamical and all three employ the quasistatic approach. Yet the Beer agents perform tasks that more closely resemble paradigmatic cognitive tasks. Phattanasri et al.’s agents were able to learn and relearn regularities linking signals to fitness-relevant features of their environment. Beer and Williams’ agents needed to maintain information about a previously observed object and then to compare this to a novel object regarding

which they were receiving and responding to information over time. In each case, the use of a quasistatic approximation played a different role in understanding how the agents’ performed their tasks. For Phattanasri’s agents, it was relevant to seeing how the agents’ optimal strategies depended both on their transitory dynamics as well as their position relative to an attractor state. For Beer and Williams’ agents, it illuminated how agents perform a task during which they need to alter their behavior to respond to constantly updating information.

This brief comparison between the three dynamical models highlights a point that should be obvious, but which is nevertheless worth making explicit. Namely, the mere fact that a system can be modeled dynamically tells one little about whether it exhibits cognition-like behavior. Given that van Gelder does not claim that the governor is a cognitive system, it is clear enough that more work needs to be done to differentiate (more) cognitive from (less, or) non-cognitive systems. But taking this point seriously requires more attention to the differences among dynamical models than is typically paid in philosophical discussions of dynamical models or explanations.

## 5 The Relationship Between Dynamical and Non-dynamical Models

The discussions in Phattanasri et al. (2007) and Beer and Williams (2015) are of philosophical interest not merely because they consider cognitively-illuminating dynamical systems, but further because they explicitly compare their dynamical representations to non-dynamical ones. While Phattanasri et al. use their dynamical models to derive a finite state machine (FSM) representation, Beer and Williams contrast their dynamical models with information theoretic ones. In this section we describe how the analysis of the dynamical systems as involving hybridity matters for understanding the relationships between these dynamical and non-dynamical representations.

Recall that Phattanasri et al. (2007) derive the FSM representation by strobing the system at various times. There is a tension in the way that they describe this representation. On the one hand, they claim that “It is important to emphasize that the extracted FSMs merely summarize the normal operation of the circuit dynamics, and are not equivalent to this dynamics” (388). On the other hand, as noted above, they talk about the FSM as being “embedded” (Phattanasri et al., 2007, 388, fig. 7) in the evolved neural circuitry and “extracted” (ibid.) from the dynamics, and of the circuits as “implementing” (291) the FSM. These comments suggest that the FSM representation is more than a mere summary of the dynamics.

So what is going on here? Should we think of the FSM representation as merely a partial and practical summary of the underlying dynamics, or as capturing a privileged pattern that is “embedded” in the neural circuitry of the minimal cognitive agents? We believe that this apparent tension can be resolved by thinking further about how the dynamical representation functions.



If one were to think of the differential equations describing a dynamical system as a pure unstructured description of the way the system evolves continuously through time, then the finite state representation will seem like a cheap reproduction of a much richer representation. Additionally, if the (so-called) underlying dynamics of the system lacked internal structure, then any discretization or decomposition would be arbitrary in the sense that it would be conceptually confused to try to carve the system at its (non-existent) joints. But the analysis of the dynamics of the system does not proceed like this.

Far from providing an unstructured description of the system’s evolutionary dynamics, Phattanasri et al. employed a quasi-static approximation in which they first created phase portraits of the attractor states of the system when its inputs are held constant and then traced the out-of-equilibrium dynamics as the system moves through its (abstract) state space in response to typical changes in the inputs. The FSM model (fig. 2c) provides an adequate representation of they dynamics insofar as it captures the transitions of the dynamical system from one attractor state towards another. The key point is that although the FSM representation introduces a discretization that did not exist in the dynamical representation, the success of this discretization can be judged based on the ability of the FSM model to capture features that are already present in the dynamical representation. Notably, the characterization of the system as involving distinct attractor states corresponding to distinct inputs and responses is an essential part of the quasi-static approximation. It is because of this feature of the modeling that it is not arbitrary to ask whether the particular finite state model delivered by strobing appropriately represents the behavior of the system, and why it is explanatory to claim that the dynamics “implement” the discrete finite state machine.

In contrast to the comparison between the FSM and dynamical representations in Phattanasri et al. (2007), in which the relationships between the representations was relatively transparent, Beer and Williams dynamical and information theoretical models do not allow for such straightforward comparisons. Unsurprisingly, there are clear links between features of their information theoretic models and features of the dynamic models employing quasi-static approximations. But there are some key differences in how the dynamic and information theoretic representations function, and it is due to these differences that they are able to play complementary but distinct roles in illuminating how the neural agents perform their tasks. We submit that careful attention to the differences in how DST and IT are used to represent their target systems will help one avoid the temptation to view one type of representation as abstracting away from the other.

We begin with an overview of how Beer and Williams use the tools of IT to model their agents. They measure the mutual information between different variables for the agents (e.g. particular neurons) and the sizes of the cues and probes. Here, the mutual information is not a relationship between the size of a particular cue in a particular trial and the activation of the neuron in that trial. Rather it depends on the different levels of the activation of the neuron corresponding to different sizes of the cue/probe across trials, and on how a

particular level of activation at a given time reduces uncertainty about the size of the cue or probe. Using the tools of IT, Beer and Williams represent the way that the information about the size of the cue is transmitted through different parts of the system through both the cue and probe stages. For instance, in the cue stage they are able to trace how the cue size information in each of the agents' internal neurons changes over time, and to thus determine where this information is maintained at the end of the cue stage.

In tracing the way that the mutual information between two variables changes over time, Beer and Williams extend information theory beyond its standard application. By providing IT approaches with a dynamic formulation, they make it easier to draw comparisons between DST approaches, which also characterize the evolution of a system. But this similarity between them potentially obscures important differences in the way that they each represent a single system. As emphasized, the mutual information between (e.g.) the cue size and a neuron's state depends on the degree to which variation in the neuron's state (at a time) tracks variation in the cue size across trials (at that time), and thus cannot be understood by reference to a trajectory in any single trial. We emphasize this point not because there is any lack of clarity in Beer and Williams' discussion, but because we believe that such nuances are easily glossed over in philosophical discussions of cognition.

As Beer and Williams describe, there are important relationships between the dynamical and informational representations of the system. For example, in representing the probe-stage dynamics, they model the mutual information between a particular neuron and a variable for the relative sizes of the cue and the probe (p. 22, fig. 10). This neuron provides the most information about relative size during the interval when the dynamic trajectories corresponding to whether the agent catches or avoids the probe are most distinct from one another (to speak somewhat imprecisely). The degree of divergence among the trajectories is an emergent property of the system corresponding to the time at which bifurcation occurs in the non-autonomous dynamics.

In Beer and Williams' discussion, IT methods serve as a proxy for computational approaches more generally. Although there is no simple way to characterize the relationships between their dynamic and information theoretic models, the relationship is certainly not what one would expect based on the philosophical literature on dynamical models. In particular, the decision to model the agent and its environment separately already exists within the quasi-static approximation employed in the dynamical models, rather than being some abstraction introduced at the computational level. The main difference we've highlighted here between DST and IT representations—that the former involve a trajectory from a single trial and the latter require cross-trial comparisons—has received no attention. This difference is important for seeing why the different methods use different bases for categorization, and thus why one cannot be understood as derived by simply abstracting away from the details of the other.

## 6 Dynamical Models and Cognition

Twenty-five years after van Gelder’s seminar paper, it is time for a more nuanced picture of dynamical models. While the Watt governor is a paradigmatic dynamical system, it is not a cognitive one, and not all claims about its dynamical model generalize to models of systems performing even “minimally cognitive” tasks. Even if van Gelder is correct that the Watt governor does not support a non-arbitrary division into discrete operational phases, it would not follow that other more complex dynamical systems fail to support such analyses. Whether a system can be fruitfully modularized is settled not by looking at whether a dynamical system describes smooth trajectories, but by careful attention to the model and the task. The dynamical models reviewed here, although simple, are significantly more complex than that of the Watt governor, and these complexities matter for determining whether they can be modeled using a more standard computational approach.

In focusing on the use of quasistatic approaches, we have made salient one way in which dynamical models are tools for representing a system. Although this point that dynamical models are themselves tools for representing a system may seem obvious, we have suggested that it gets lost in the setup of current debates. While philosophers defending a particular computational or mechanistic model need to be explicit about how the model divides up the system, the representational choices underlying dynamical models are typically left implicit. Yet dynamical models do not, in general, provide unstructured descriptions of a system’s temporal evolution, and careful attention to the devices by which time is modeled yields insights into the conditions under which the models apply.

What, ultimately, is the relationship between dynamical and non-dynamical representations of a cognitive system? There is no general answer to this question. Dynamical and non-computational models are not mutually exclusive, and very little can be inferred from the mere fact that a system can be modeled in one framework or the other. This, in fact, requires philosophers to pay *more* attention to the formal features of particular models, since whether a system ought to be modeled computationally can only be resolved by extended attention to how the agent’s dynamics enable it to perform its task. Dynamical models of cognition will not replace computational ones, but promise a deeper understanding of how computational systems work.

## References

- Bechtel W (1998) Representations and cognitive explanations: Assessing the dynamicist’s challenge in cognitive science. *Cognitive Science* 22(3):295–317
- Beer RD, Williams PL (2015) Information processing and dynamics in minimally cognitive agents. *Cognitive science* 39(1):1–38
- Bokulich A (2011) How scientific models can explain. *Synthese* 180(1):33–45

- Cat J (2005) Modeling cracks and cracking models: structures, mechanisms, boundary conditions, constraints, inconsistencies and the proper domains of natural laws. *Synthese* 146(3):447–487
- Chemero A, Silberstein M (2008) After the philosophy of mind: Replacing scholasticism with science. *Philosophy of science* 75(1):1–27
- Clark A (1998) Embodiment and the philosophy of mind. *Royal Institute of Philosophy Supplements* 43:35–51
- Dash D, Druzdzel M (2001) Caveats for causal reasoning with equilibrium models. In: *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, Springer, pp 192–203
- Eliasmith C (2010) How we ought to describe computation in the brain. *Studies in History and Philosophy of Science Part A* 41(3):313–320
- Fodor JA (1981) *Representations: Philosophical essays on the foundations of cognitive science*. Mit Press Cambridge, MA
- van Gelder T (1995) What might cognition be, if not computation? *The Journal of Philosophy* 92(7):345–381
- Green S, Batterman R (2017) Biology meets physics: Reductionism and multi-scale modeling of morphogenesis. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 61:20–34
- Grush R (1997) Yet another design for a brain? review of port and van gelder. *Mind as motion Philosophical Psychology* 10:233–242
- Haken H, Kelso JS, Bunz H (1985) A theoretical model of phase transitions in human hand movements. *Biological cybernetics* 51(5):347–356
- Hempel CG, Oppenheim P (1948) Studies in the logic of explanation. *Philosophy of science* 15(2):135–175
- Iwasaki Y, Simon HA (1994) Causality and model abstraction. *Artificial intelligence* 67(1):143–194
- Kaplan DM (2015) Moving parts: the natural alliance between dynamical and mechanistic modeling approaches. *Biology & Philosophy* 30(6):757–786
- Kaplan DM, Craver CF (2011) The explanatory force of dynamical and mathematical models in neuroscience: A mechanistic perspective. *Philosophy of science* 78(4):601–627
- Phattanasri P, Chiel HJ, Beer RD (2007) The dynamics of associative learning in evolved model circuits. *Adaptive behavior* 15(4):377–396

- Port RF, van Gelder T (1995) *Mind as motion: Explorations in the dynamics of cognition*. MIT press
- Putnam H (1960) *Minds and machines*. In: Hook S (ed) *Dimensions of Mind: A Symposium*, New York University Press, pp 148–179
- Putnam H (1967) *The nature of mental states*. originally published as psychological predicates. *Art, Mind, and Religion*
- Rescher N, Simon HA (1966) *Cause and counterfactual*. *Philosophy of Science* 33(4):323–340
- Thelen E, Smith LB (1996) *A dynamic systems approach to the development of cognition and action*. MIT press
- Weinberger N (2019) *Reintroducing dynamics into static causal models*. In: Kleinberg S (ed) *Time and Causality Across the Sciences*, Cambridge University Press, pp 187–211
- Wheeler M (2005) *Reconstructing the cognitive world: The next step*. MIT press
- Wilkenfeld DA (2014) *Functional explaining: A new approach to the philosophy of explanation*. *Synthese* 191(14):3367–3391
- Wilson M (2017) *Physics Avoidance: and other essays in conceptual strategy*. Oxford University Press
- Zednik C (2011) *The nature of dynamical explanation*. *Philosophy of Science* 78(2):238–263