Journal of Philosophy, Inc.

What Might Cognition Be, If Not Computation?
Author(s): Tim Van Gelder
Source: *The Journal of Philosophy*, Vol. 92, No. 7 (Jul., 1995), pp. 345-381
Published by: Journal of Philosophy, Inc.
Stable URL: http://www.jstor.org/stable/2941061
Accessed: 30/10/2008 13:51

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at http://www.jstor.org/page/info/about/policies/terms.jsp. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at http://www.jstor.org/action/showPublisher?publisherCode=jphil.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit organization founded in 1995 to build trusted digital archives for scholarship. We work with the scholarly community to preserve their work and the materials they rely upon, and to build a common research platform that promotes the discovery and use of these resources. For more information about JSTOR, please contact support@jstor.org.

Journal of Philosophy, Inc. is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Philosophy.*

http://www.jstor.org

THE JOURNAL OF PHILOSOPHY
VOLUME XCI, NO. 7, JULY 1995

WHAT MIGHT COGNITION BE, IF NOT COMPUTATION?*

What is cognition? Contemporary orthodoxy maintains that it is computation: the mind is a special kind of computer, and cognitive processes are the rule-governed manipulation of internal symbolic representations. This broad idea has dominated the philosophy and the rhetoric of cognitive science—and even, to a large extent, its practice—ever since the field emerged from the postwar cybernetic melee. It has provided the general framework for much of the most well-developed and insightful research into the nature of mental operation. Yet, over the last decade or more, the computational vision has lost much of its lustre. Although work within it continues apace, a variety of difficulties and limitations have become increasingly apparent, and researchers across cognitive science and related disciplines have been casting around for other ways to understand cognitive processes. Partly as a result, there are now many research programs which, one way or another, stand opposed to the traditional computational approach; these include connectionism, neurocomputational approaches, ecological psychology, situated robotics, synergetics, and artificial life.

These approaches appear to offer a variety of differing and even conflicting conceptions of the nature of cognition. It is therefore an appropriate time to step back and reconsider the question: What general arguments are there in favor of the idea that cognitive processes must be specifically *computational* in nature? In order prop-

---

* Criticism and advice from numerous people helped improve this paper, but special acknowledgement is due to Robert Port, John Haugeland, and James Townsend. Audiences at the University of Illinois/Chicago, the New Mexico State University, Indiana University, the Australian National University, the University of New South Wales, Princeton University, Lehigh University, and the University of Skövde were suitably and helpfully critical of earlier versions.

0022-362X/95/9207/345-81

erly to address this question, however, we must first address another: What are the alternatives? What *could* cognition be, if it were *not* computation of some form or other?

There are at least two reasons why this second question is important. First, arguments in favor of some broad hypothesis are rarely, if ever, completely general. They tend to be arguments not for *A* alone, but rather in favor of *A* as opposed to *B*, and such arguments often fail to support *A* as opposed to *C*. For example, one of the most powerful early considerations raised in favor of the computational conception of cognition was the idea that intelligent behavior requires sophisticated internal representations. While this clearly supported the computational conception against a behaviorism which eschewed such resources, however, it was no use against a connectionism which helped itself to internal representations, though rather different in kind than the standard symbolic variety.

The second reason we need to ask what alternatives there may be is that one of the most influential arguments in favor of the computational view is the claim that there is simply no alternative. This is sometimes known as the *"what else could it be?"* argument.[1] As Allen Newell[2] recently put it:

> ...although a small chance exists that we will see a new paradigm emerge for mind, it seems unlikely to me. Basically, there do not seem to be any viable alternatives. This position is not surprising. In lots of sciences we end up where there are no major alternatives around to the particular theories we have. Then, all the interesting kinds of scientific action occur inside the major view. It seems to me that we are getting rather close to that situation with respect to the computational theory of mind (*ibid.*, p. 56).

This paper describes a viable alternative. Rather than computers, cognitive systems may be dynamical systems; rather than computation, cognitive processes may be state-space evolution within these very different kinds of systems. It thus disarms the "what else could it be?" argument, and advances the broader project of evaluating competing hypotheses concerning the nature of cognition. Note that achieving these goals does not require decisively establishing that the dynamical hypothesis is true. That would require considerably more space than is available here, and to attempt it now would be hopelessly premature anyway. All that must be done is to describe

---

[1] This title may have been first used in print by John Haugeland in "The Nature and Plausibility of Cognitivism," *Behavioral and Brain Sciences*, I (1978): 215–26.

[2] "Are There Alternatives?" in W. Sieg, ed., *Acting and Reflecting* (Boston: Kluwer, 1990).

and motivate the dynamical conception sufficiently to show that it does in fact amount to an alternative conception of cognition, and one which is currently viable, as far as we can now tell.

A fruitful way to present the dynamical conception is to begin with an unusual detour, via the early industrial revolution in England, circa 1788.

I. THE GOVERNING PROBLEM

A central engineering challenge for the industrial revolution was to find a source of power that was reliable, smooth, and uniform. In the latter half of the eighteenth century, this had become the problem of translating the oscillating action of the steam piston into the rotating motion of a flywheel. In one of history's most significant technological achievements, Scottish engineer James Watt designed and patented a gearing system for a rotative engine. Steam power was no longer limited to pumping; it could be applied to any machinery that could be driven by a flywheel. The cotton industry was particularly eager to replace its horses and water wheels with the new engines. High-quality spinning and weaving required, however, that the source of power be highly uniform, that is, there should be little or no variation in the speed of revolution of the main driving flywheel. This is a problem, since the speed of the flywheel is affected both by the pressure of the steam from the boilers, and by the total workload being placed on the engine, and these are constantly fluctuating.

It was clear enough how the speed of the flywheel had to be regulated. In the pipe carrying steam from the boiler to the piston there was a throttle valve. The pressure in the piston, and so the speed of the wheel, could be adjusted by turning this valve. To keep engine speed uniform, the throttle valve would have to be turned, at just the right time and by just the right amount, to cope with changes in boiler pressure and workload. How was this to be done? The most obvious solution was to employ a human mechanic to turn the valve as necessary. This had a number of drawbacks, however: mechanics required wages, and were often unable to react sufficiently swiftly and accurately. The industrial revolution thus confronted a second engineering challenge: design a device which can automatically adjust the throttle valve so as to maintain uniform speed of the flywheel despite changes in steam pressure or workload. Such a device is known as a *governor.*

Difficult engineering problems are often best approached by breaking the overall task down into simpler subtasks, continuing the process of decomposition until one can see how to construct devices that can directly implement the various component tasks. In the case of the governing problem, the relevant decomposition seems clear. A change need only be made to the throttle valve if the flywheel is not currently running at the correct speed. Therefore, the first subtask must be to measure the speed of the wheel, and the second subtask must be to calculate whether there is any discrepancy between the desired speed and the actual speed. If there is no discrepancy, no change is needed, for the moment at least. If there is a discrepancy, then the governor must determine by how much the throttle valve should be adjusted to bring the speed of the wheel to the desired level. This will depend, of course, on the current steam pressure, and so the governor must measure the current steam pressure and then on that basis calculate how much to adjust the valve. Finally, of course, the valve must be adjusted. This overall sequence of subtasks must be carried out as often as necessary to keep the speed of the wheel sufficiently close to the desired speed.

A device that can solve the governing problem would have to carry out these various subtasks repeatedly in the correct order, and so we can think of it as obeying the following algorithm:

1. Measure the speed of the flywheel.
2. Compare the actual speed against the desired speed.
3. If there is no discrepancy, return to step 1. Otherwise,
   a. measure the current steam pressure;
   b. calculate the desired alteration in steam pressure;
   c. calculate the necessary throttle valve adjustment.
4. Make the throttle valve adjustment.

Return to step 1.

There must be some physical device capable of actually carrying out each of these subtasks, and so we can think of the governor as incorporating a tachometer (for measuring the speed of the wheel); a device for calculating the speed discrepancy; a steam pressure meter; a device for calculating the throttle valve adjustment; a throttle valve adjuster; and some kind of central executive to handle sequencing of operations. This conceptual breakdown of the components of the governor may even correspond to its actual breakdown; that is, each of these components may be implemented by a distinct, dedicated physical device. The engineering problem would then reduce to the (presumably much simpler) problem of constructing the various components and hooking them together so that the whole system functions in a coherent fashion.

Now, as obvious as this approach now seems, it was not the way the governing problem was actually solved. For one thing, it presupposes devices that can swiftly perform some quite complex calculations, and although some simple calculating devices had been invented in the seventeenth century, there was certainly nothing available in the late eighteenth century that could have met the demands of a practical governor.

The real solution, adapted by Watt from existing windmill technology, was much more direct and elegant. It consisted of a vertical spindle geared into the main flywheel so that it rotated at a speed directly dependent upon that of the flywheel itself (see figure 1). Attached to the spindle by hinges were two arms, and on the end of each arm was a metal ball. As the spindle turned, centrifugal force drove the balls outward and hence upward. By a clever arrangement, this arm motion was linked directly to the throttle valve. The result was that as the speed of the main wheel increased, the arms raised, closing the valve and restricting the flow of steam; as the speed decreased, the arms fell, opening the valve and allowing more steam to flow. The engine adopted a constant speed, maintained with extraordinary swiftness and smoothness in the presence of large fluctuations in pressure and load.

It is worth emphasizing how remarkably well the centrifugal governor actually performed its task. This device was not just an engineer-

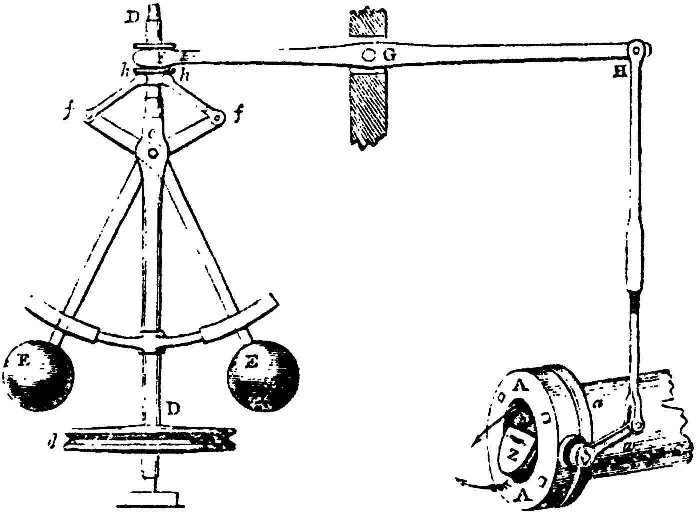

Figure 1[3]

---

[3] The Watt centrifugal governor for controlling the speed of a steam engine—from J. Farey, *A Treatise on the Steam Engine: Historical, Practical, and Descriptive* (London: Longman, Rees, Orme, Brown, and Green, 1827).

ing hack employed because computer technology was unavailable. *Scientific American* claimed in 1858 that an American variant of the basic centrifugal governor, "if not absolutely perfect in its action, is so nearly so, as to leave in our opinion nothing further to be desired."

But why should any of this be of any interest in the philosophy of cognitive science? The answer may become apparent as we examine a little more closely some of the differences between the two governors.

II. TWO KINDS OF GOVERNORS

The two governors described in the previous section are patently different in construction, yet they both solve the same control problem, and we can assume (for purposes of this discussion) that they both solve it sufficiently well. Does it follow that, deep down, they are really the same kind of device, despite superficial differences in construction? Or are they deeply different, despite their similarity in overt performance?

It is natural to think of the first governor as a computational device; one which, as part of its operation computes some result, namely, the desired change in throttle valve angle. Closer attention reveals that there is in fact a complex group of properties here, a group whose elements are worth teasing apart.

Perhaps the most central of the computational governor's distinctive properties is its dependence on representation. Every aspect of its operation, as outlined above, deals with representations in some manner or other. The very first thing it does is measure its environment (the engine) to obtain a symbolic representation of current engine speed. It then performs a series of operations on this and other representations, resulting in an output representation, a symbolic specification of the alteration to be made in the throttle valve; this representation then causes the valve adjusting mechanism to make the corresponding change. This is why it is appropriately described as computational (now in a somewhat narrower sense): it literally computes the desired change in throttle valve by manipulating symbols according to a schedule of rules. Those symbols, in the context of the device and its situation, have meaning, and the success of the governor in its task is owed to its symbol manipulations being in systematic accord with those meanings. The manipulations are discrete operations which necessarily occur in a determinate sequence; for example, the appropriate change in the throttle valve can only be calculated after the discrepancy between current and desired speeds has been calculated. At the highest level, the whole device operates in a cyclic fashion: it first measures (or "perceives") its environment; it then internally computes an appropriate change in throttle valve; it then effects this change ("acts" on its environment). After the change has been made and given time to affect engine speed, the governor runs through whole the cycle again...and again.... Finally, notice that the governor is homuncular in construction. Homuncularity is a special kind of breakdown of a system into parts or components, each of which is responsible for a particular subtask. Homuncular components are ones that, like departments or committees within bureaucracies, interact by communication (that is, by passing meaningful messages). Obviously, the representational and computational nature of the governor is essential to its homuncular construction: if the system as a whole did not operate by manipulating representations, it would not be possible for its components to interact by communication.

These properties—representation, computation, sequential and cyclic operation, and homuncularity—form a mutually interdependent cluster; a device with any one of them will standardly possess others. Now, the Watt centrifugal governor does not exhibit this cluster of properties as a whole, nor any one of them individually. As obvious as this may seem, it deserves a little detailed discussion and argument, since it often meets resistance, and some useful insights can be gained along the way.

Since manipulable representations lie at the heart of the computational picture, the nonrepresentational nature of the centrifugal governor is a good place to start. There is a common and initially quite attractive intuition to the effect that the angle at which the arms are swinging is a representation of the current speed of the engine, and it is because the arms are related in this way to engine speed that the governor is able to control that speed. This intuition is misleading, however; arm angle and engine speed are of course intimately related, but the relationship is not representational. There are a number of powerful arguments favoring this conclusion. They are not based on any unduly restrictive definition of the notion of representation; they go through on pretty much any reasonable characterization, based around a core idea of some state of a system which, by 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.[4]

---

[4] This broad characterization is adapted from Haugeland, "Representational Genera," in W. Ramsey, S.P. Stich, D.E. Rumelhart, eds., *Philosophy and Connectionist Theory* (Hillsdale, NJ: Erlbaum, 1991), pp. 61–89.

A useful criterion of representation—a reliable way of telling whether a system contains them or not—is to ask whether there is any explanatory utility in describing the system in representational terms. If you really can make substantially more sense of how a system works by concretely describing various identifiable parts or aspects of it as representations in the above sense, that is the best evidence you could have that the system really does contain representations. Conversely, if describing the system as representational lets you explain nothing over and above what you could explain before, why on earth suppose it to be so? Note that very often representational descriptions do yield substantial explanatory benefits. This is certainly true for pocket calculators, and mainstream cognitive science is premised on the idea that humans and animals are like that as well. A noteworthy fact about standard explanations of how the centrifugal governor works is, however, that they never talk about representations. This was true for the informal description given above, which apparently suffices for most readers; more importantly, it has been true of the much more detailed descriptions offered by those who have actually been in the business of constructing centrifugal governors or analyzing their behavior. Thus, for example, a mechanics manual for construction of governors from the middle of last century, Maxwell's original dynamical analysis (see below), and contemporary mathematical treatments all describe the arm angle and its role in the operation of the governor in nonrepresentational terms. The reason, one might reasonably conclude, is that the governor contains no representations.

The temptation to treat the arm angle as a representation comes from the informal observation that there is some kind of correlation between arm angle and engine speed; when the engine rotates at a certain speed, the arms will swing at a given angle. Now, supposing for the moment that this is an appropriate way to describe their relationship, it would not follow that the arm angle is a representation. One of the few points of general agreement in the philosophy of cognitive science is that mere correlation does not make something a representation. Virtually everything is correlated, fortuitously or otherwise, with something else; to describe every correlation as representation is to trivialize representation. For the arm angle to count, in the context of the governing system alone, as a representation, we would have to be told what else about it justifies the claim that it is a representation.

But to talk of some kind of correlation between arm angle and engine speed is grossly inadequate, and once this is properly understood, there is simply no incentive to search for this extra ingredient.

For a start, notice that the correlation at issue only obtains when the total system has reached its stable equilibrium point, and is immediately disturbed whenever there is some sudden change in, for example, the workload on the engine. At such times, the speed of the engine quickly drops for a short period, while the angle of the arms adjusts only at the relatively slow pace dictated by gravitational acceleration. Yet, even as the arms are falling, more steam is entering the piston, and hence the device is already working; indeed, these are exactly the times when it is most crucial that the governor work effectively. Consequently, no simple correlation between arm angle and engine speed can be the basis of the operation of the governor.

The fourth and deepest reason for supposing that the centrifugal governor is not representational is that, when we fully understand the relationship between engine speed and arm angle, we see that the notion of representation is just the wrong sort of conceptual tool to apply. There is no doubt that at all times the arm angle is in some interesting way related to the speed of the engine. This is the insight which leads people to suppose that the arm angle is a representation. Yet appropriately close examination of this dependence shows exactly why the relationship cannot be one of representation. For notice that, because the arms are directly linked to the throttle valve, the angle of the arms is at all times determining the amount of steam entering the piston, and hence at all times the speed of the engine depends in some interesting way on the angle of the arms. Thus, arm angle and engine speed are at all times both determined by, and determining, each other's behavior. As we shall see below, there is nothing mysterious about this relationship; it is quite amenable to mathematical description. Yet it is much more subtle and complex than the standard concept of representation can handle, even when construed as broadly as is done here. In order to describe the relationship between arm angle and engine speed, we need a more powerful conceptual framework than mere talk of representations. That framework is the mathematical language of dynamics, and in that language, the two quantities are said to be coupled. The real problem with describing the governor as a representational device, then, is that the relation of representing—something standing in for some other state of affairs—is too simple to capture the actual interaction between the governor and the engine.

If the centrifugal governor is not representational, then it cannot be computational, at least in the specific sense that its processing cannot be a matter of the rule-governed manipulation of symbolic representations. Its noncomputational nature can also be established another way. Not 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 nonarbitrarily dividing its changes over time into distinct manipulatings, and no point in trying to do so. From this, it follows that the centrifugal governor is not sequential and not cyclic in its operation in anything like the manner of the computational governor. Since there are no distinct processing steps, there can be no sequence in which those steps occur. There is never any one operation that must occur before another one can take place. Consequently, there is nothing cyclical about its operation. The device has, to be sure, an "input" end (where the spindle is driven by the engine) and an "output" end (the connection to the throttle valve). But the centrifugal governor does not follow a cycle where it first takes a measurement, then computes a throttle valve change, then makes that adjustment, then takes a measurement, and so on. Rather, input, internal activity, and output are all happening continuously and at the very same time, much as a radio is producing music at the very same time as its antenna is receiving signals.

The fact that the centrifugal governor is not sequential or cyclic in any respect points to yet another deep difference between the two kinds of governor. There is an important sense in which time does not matter in the operation of the computational governor. There is, of course, the minimal constraint that the device must control the engine speed adequately, and so individual operations within the device must be sufficiently fast. There is also the constraint that internal operations must happen in the right sequence. Beyond these, however, there is nothing that dictates when each internal operation takes place, how long it takes to carry it out, and how long elapses between each operation. There are only pragmatic implementation considerations: which algorithms to use, what kind of hardware to use to run the algorithms, and so forth. The timing of the internal operations is thus essentially arbitrary relative to that of any wider course of events. It is as if the wheel said to the governing system: "Go away and figure out how much to change the valve to keep me spinning at 100 rpm. I don't care how you do it, how many steps you take, or how long you take over each step, as long as you report back within (say) 10 milliseconds."

In the centrifugal governor, by contrast, there is simply nothing that is temporally unconstrained in this way. There are no occurrences whose timing is arbitrary relative to the operation of the engine. All behavior in the centrifugal governor happens in the very same real time frame as change in the speed of the flywheel. We can sum up the point this way: the two kinds of governor differ fundamentally in their temporality, and the temporality of the centrifugal governor is essentially that of the engine itself.

Finally, it need hardly be labored that the centrifugal governor is not a homuncular system. It has parts, to be sure, and its overall behavior is the direct result of the organized interaction of those parts. The difference is that those parts are not modules interacting by communication; they are not like little bureaucratic agents passing representations among themselves as the system achieves the overall task.

III. CONCEPTUAL FRAMEWORKS

In the previous section, I argued that the differences in nature between the two governors run much more deeply than the obvious differences in mechanical construction. Not surprisingly, these differences in nature are reflected in the kind of conceptual tools that we must bring to bear if we wish to understand the operation of these devices. That is, the two different governors require very different conceptual frameworks in order to understand how it is that they function as governors, that is, how they manage to control their environment.

In the case of the computational governor, the behavior is captured in all relevant detail by an algorithm, and the general conceptual framework we are bringing to bear is that of mainstream computer science. Computer scientists are typically concerned with what you can achieve by stringing together, in an appropriate order, some set of basic operations: either how best to string them together to achieve some particular goal (programming, theory of algorithms), or what is achievable in principle in this manner (computation theory). So we understand the computational governor as a device capable of carrying out some set of basic operations (measurings, subtractings, etc.), and whose sophisticated overall behavior results from nothing more than the complex sequencing of these basic operations. Note that there is a direct correspondence between elements of the governor (the basic processing steps it goes through) and elements of the algorithm which describes its operation (the basic instructions).

The Watt centrifugal governor, by contrast, cannot be understood this way at all. There is nothing in that device for any algorithm to lock onto. Very different conceptual tools have always been applied to this device. The terms in which it was described above, and indeed by Watt and his peers, were straightforwardly mechanical: rotations, spindles, levers, displacements, forces. Last century, more precise and powerful descriptions became available, but these also have nothing to do with computer science. In 1868, the physicist James Clerk Maxwell[5] made a pioneering extension of the mathematical tools of dynamics to regulating and governing devices. The general approach he established has been standard ever since. Though familiar to physicists and control engineers, it is less so to most cognitive scientists and philosophers of mind, and hence is worth describing in a little detail.

The key feature of the governor's behavior is the angle at which the arms are hanging, for this angle determines how much the throttle valve is opened or closed. Therefore, in order to understand the behavior of the governor, we need to understand the basic principles governing how arm angle changes over time. Obviously, the arm angle depends on the speed of the engine; hence we need to understand change in arm angle as a function of engine speed. If we suppose for the moment that the link between the governor and the throttle valve is disconnected, then this change is given by the differential equation:

$$\frac{d^2\theta}{dt^2} = (n\omega)^2 \cos\theta \sin\theta - \frac{g}{l}\sin\theta - r\frac{d\theta}{dt}$$

where θ is the angle of arms, n is a gearing constant, ω is the speed of engine, g is a constant for gravity, l is the length of the arms, and r is a constant of friction at hinges.[6] This nonlinear, second-order differential equation tells us the instantaneous acceleration in arm angle, as a function of what the current arm angle happens to be (designated by the state variable θ), how fast arm angle is currently changing (the derivative of θ with respect to time, $d\theta/dt$) and the current engine speed (ω). In other words, the equation tells us how change in arm angle is changing, depending on the current arm angle, the way it is changing already, and the engine speed. Note that in the system defined by this equation, change over time occurs only in arm angle θ (and its derivatives). The other quantities (ω, n, g, l, and r) are assumed to stay fixed, and are called parameters. The particular values at which the parameters are fixed determine the precise shape of the change in θ. For this reason, the parameter settings are said to fix the dynamics of the system.

This differential equation is perfectly general and highly succinct: it is a way of describing how the governor behaves for any arm angle and engine speed. This generality and succinctness comes at a price, however. If we happen to know what the current arm angle is, how fast it is changing, and what the engine speed is, then from this

---

[5] "On Governors," *Proceedings of the Royal Society*, XVI (1868): 270-83.

[6] Edward Beltrami, *Mathematics for Dynamical Modeling* (Boston: Academic, 1987), p. 163.

equation all we can figure out is the current instantaneous acceleration. If we want to know at what angle the arms will be in a half-second, for example, we need to find a solution to the general equation—that is, an equation that tells us what values θ takes as a function of time, which satisfies the differential equation. There are any number of such solutions, corresponding to all the different behavioral trajectories that the governor might exhibit, but these solutions often have important general properties in common; thus, as long as the parameters stay within certain bounds, the arms will always eventually settle into a particular angle of equilibrium for that engine speed; that angle is known as a *point attractor.*

Thus far I have been discussing the governor without taking into account its effect on the engine, and thereby indirectly on itself. Here, the situation gets a little more complicated, but the same mathematical tools apply. Suppose we think of the steam engine itself as a dynamical system governed by a set of differential equations, one of which gives us some derivative of engine speed as a function of current engine speed and a number of other variables and parameters:

$$\frac{d^n\omega}{dt^n} = F(\omega, \ldots, \tau, \ldots)$$

One of these parameters is the current setting of the throttle valve, τ, which depends directly on the governor arm angle θ. We can thus think of θ as a parameter of the engine system, just as engine speed ω is a parameter of the governor system. (Alternatively, we can think of the governor and steam engine as comprising a single dynamical system in which both arm angle and engine speed are state variables.) This relationship, known as *coupling,* is particularly interesting and subtle. Changing a parameter of a dynamical system changes its total dynamics (that is, the way its state variables change their values depending on their current values, across the full range of values they may take). Thus, any change in engine speed, no matter how small, changes not the state of the governor directly, but rather the way the state of the governor *changes,* and any change in arm angle changes the way the state of the engine changes. Again, however, the overall system (coupled engine and governor) settles quickly into a point attractor, that is, engine speed and arm angle remain constant. Indeed, the remarkable thing about this coupled system is that under a wide variety of conditions it always settles swiftly into states at which the engine is running at a particular speed. This is of course exactly what is wanted: coupling the governor to the engine results in the engine running at a constant speed.

In this discussion, two very broad, closely related sets of conceptual resources have (in a modest way) been brought into play. The first is dynamical modeling, that branch of applied mathematics which attempts to describe change in real-world systems by describing the states of the system numerically and then writing equations that capture how these numerical states change over time. The second set of resources is dynamical systems theory, the general study of dynamical systems considered as abstract mathematical structures. Roughly speaking, dynamical modeling attempts to understand natural phenomena as the behavior of real-world realizations of abstract dynamical systems, whereas dynamical systems theory studies the abstract systems themselves. There is no sharp distinction between these two sets of resources, and for our purposes they can be lumped together under the general heading of dynamics.

IV. MORALS

This discussion of the governing task suggests a number of closely related lessons for cognitive science:

(1) Various different kinds of systems, fundamentally different in nature and requiring very different conceptual tools for their understanding, can subserve sophisticated tasks—including interacting with a changing environment—which may initially appear to demand that the system have knowledge of, and reason about, its environment. The governing problem is one simple example of such a task; it can be solved either by a computational system or by a noncomputational dynamical system, the Watt centrifugal governor.

(2) In any given case, our sense that a specific cognitive task *must* be subserved by a (generically) computational system may be due to deceptively compelling preconceptions about how systems solving complex tasks must work. Many people are oblivious to the possibility of a noncomputational, dynamical solution to the governing problem, and so all-too-readily assume that it must be solved in a computational manner. Likewise, it may be that the basically computational shape of most mainstream models of cognition results not so much from the nature of cognition itself as it does from the shape of the conceptual equipment that cognitive scientists typically bring to bear in studying cognition.

(3) Cognitive systems may in fact be *dynamical* systems, and cognition the behavior of some (noncomputational) dynamical system. Perhaps, that is, cognitive systems are more relevantly similar to the centrifugal governor than they are similar either to the computational governor, or to that more famous exemplar of the broad category of computational systems, the Turing machine.

In what follows, the first and third of these points will be elaborated in just enough detail to substantiate the basic claim of this paper, that there is in fact a currently viable alternative to the computational conception of cognition. As a first step toward doing that, however, I shall briefly describe an example of dynamical research in cognitive science, in order to provide what might seem to be no more than rank speculation with a little healthy flesh.

V. AN EXAMPLE OF DYNAMICAL RESEARCH

Consider the process of coming to make a decision between a variety of options, each of which has attractions and drawbacks. This is surely a high-level cognitive task, if anything is. Psychologists have done endless experimental studies determining how people choose, and produced many mathematical models attempting to describe and explain their choice behavior. The dominant approach in modeling stems from the classic expected-utility theory and statistical decision theory as originally developed by John von Neumann and Oskar Morgenstern. The basic idea here is that an agent makes a decision by selecting the option that has the highest expected utility, which is calculated in turn by combining some formal measure of the utility of any given possible outcome with the probability that it will eventuate if the option is chosen. Much of the work within the classical framework is mathematically elegant and provides a useful description of optimal reasoning strategies. As an account of the actual decisions people reach, however, classical utility theory is seriously flawed; human subjects typically deviate from its recommendations in a variety of ways. As a result, many theories incorporating variations on the classical core have been developed, typically relaxing certain of its standard assumptions, with varying degrees of success in matching actual human choice behavior. Nevertheless, virtually all such theories remain subject to some further drawbacks:

(1) They do not incorporate any account of the underlying motivations that give rise to the utility that an object or outcome holds at a given time.

(2) They conceive of the utilities themselves as static values, and can offer no good account of how and why they might change over time, and why preferences are often inconsistent and inconstant.

(3) They offer no serious account of the deliberation process, with its attendant vacillations, inconsistencies, and distress; and they have

nothing to say about the relationships that have been uncovered between time spent deliberating and the choices eventually made.

Curiously, these drawbacks appear to have a common theme; they all concern, one way or another, *temporal* aspects of decision making. It is worth asking whether they arise because of some deep structural feature inherent in the whole framework which conceptualizes decision-making behavior in terms of calculating expected utilities.

Notice that utility-theory based accounts of human decision making ("utility theories") are deeply akin to the computational solution to the governing task. That is, if we take such accounts as not just describing the outcome of decision-making behavior, but also as a guide to the structures and processes that underlie such behavior,[7] then there are basic structural similarities to the computational governor. Thus, utility theories are straightforwardly computational; they are based on static representations of options, utilities, probabilities, and so on, and processing is the algorithmically specifiable internal manipulation of these representations to obtain a final representation of the choice to be made. Consequently, utility theories are strictly sequential; they presuppose some initial temporal stage at which the relevant information about options, likelihoods, and so on, is acquired; a second stage in which expected utilities are calculated; and a third stage at which the choice is effected in actual behavior. And, like the computational governor, they are essentially atemporal; there are no inherent constraints on the timing of the various internal operations with respect to each other or change in the environment.

What we have, in other words, is a model of human cognition which, on one hand, instantiates the same deep structure as the computational governor, and on the other, seems structurally incapable of accounting for certain essentially temporal dimensions of decision-making behavior. At this stage, we might ask: What kind of model of decision-making behavior we would get if, rather, we took the *centrifugal* governor as a prototype? It would be a model with a relatively small number of continuous variables influencing each other in real time. It would be governed by nonlinear differential equations. And it would be a model in which the agent and the choice environment, like the governor and the engine, are tightly interlocked.

---

[7] See, for example, J.W. Payne, J.R. Bettman, and E.J. Johnson, "Adaptive Strategy Selection in Decision Making," *Journal of Experimental Psychology: Learning, Memory, Cognition*, XIV (1988): 534–52.

It would, in short, be rather like the *motivational oscillatory theory* (MOT) modeling framework described by mathematical psychologist James Townsend.[8] MOT enables modeling of various qualitative properties of the kind of cyclical behaviors that occur when circumstances offer the possibility of satiation of desires arising from more or less permanent motivations; an obvious example is regular eating in response to recurrent natural hunger. It is built around the idea that in such situations, your underlying motivation, transitory desires with regard to the object, distance from the object, and consumption of it are continuously evolving and affecting each other in real time; for example, if your desire for food is high and you are far from it, you will move toward it (that is, z changes), which influences your satiation and so your desire. The framework thus includes variables for the current state of motivation, satiation, preference, and action (movement), and a set of differential equations describe how these variables change over time as a function of the current state of the system.[9]

---

[8] See "A Neuroconnectionistic Formulation of Dynamic Decision Field Theory," in D. Vickers and P.L. Smith, eds., *Human Information Processing: Measures, Mechanisms, and Models* (Amsterdam: North Holland, 1988); and "Don't Be Fazed by PHASER: Beginning Exploration of a Cyclical Motivational System," *Behavior Research Methods, Instruments and Computers*, xxiv (1992): 219–27.

[9] The equations, with rough and partial translations into English, are:

$$\frac{dx}{dt} = M - m - c$$

(The change in motivation depends on how the current levels of motivation and of consumption compare with some standard level of motivation, M.)

$$\frac{dx}{dt} = \left[\frac{1}{z_1^2 + z_2^2 + a} + 1\right] \cdot m$$

(The change in one's *preference* for the goal will depend on current motivation and one's distance from the object of preference.)

$$\frac{dc}{dt} = (x + C - c) \cdot \left[\frac{b}{z_1^2 + z_2^2 + r}\right]$$

(The change in consumption will depend on the level of preference, the level of consumption, and the distance from the object of preference.)

$$\frac{dz_1}{dt} = -x \cdot z_1 \qquad \frac{dz_2}{dt} = -x \cdot z_2$$

(How one moves toward or away from the object depends on one's current level of preference for the object.) See "Don't Be Fazed by PHASER" for an accessible and graphic introduction to the behaviors defined by these equations.

MOT stands to utility theories in much the same relation as the centrifugal governor does to the computational governor. In MOT, cognition is not the manipulation of symbols, but rather state-space evolution in a dynamical system. MOT models produce behavior which, if one squints while looking at it, seems like decision making—after all, the agent will make the move which offers the most reward, which in this case means moving toward food if sufficiently hungry. But this is decision making without decisions, so to speak, for there never are in the model any discrete internal occurrences that one could characterize as decisions. In this approach, decision making is better thought of as the behavior of an agent under the influence of the pushes and pulls that emanate from desirable outcomes, undesirable outcomes, and internal desires and motivations; in a quasi-gravitational way, these forces act on the agent with strength varying as a function of distance.

The MOT modeling framework is a special case of a more general (and rather more complex) dynamical framework which Townsend and Jerome Busemeyer[10] call "decision field theory." That framework allows faithful modeling of a wide range of behaviors more easily recognizable as decision making as studied within the traditional research paradigm; indeed, their claim is that decision field theory "covers a broader range of phenomena in greater detail" than classical utility theories, and even goes beyond them by explaining in a natural way several important paradoxes of decision making, such as the so-called "common consequence effect" and "common ratio effect." The important point for immediate purposes, however, is that the general decision field theory works on the same fundamental dynamical principles as MOT. There is thus no question that at least certain aspects of human high-level cognitive functioning can be modeled effectively using dynamical systems of the kind that can be highlighted by reference to the centrifugal governor.

Thus far, all I have done is to use the governing problem as a means of exploring some of the deep differences between computational and noncomputational solutions to complex tasks, drawn out some suggestive implications for cognitive science, and used the Busemeyer and Townsend work to illustrate the claim that high-level cognitive processes can in fact be modeled using noncomputa-

---

[10] "Decision Field Theory: A Dynamic-Cognitive Approach to Decision Making in an Uncertain Environment," *Psychological Review*, C (1993): 432–59; an accessible overview is given in "Dynamic Representation of Decision Making," in R. Port and myself, eds., *Mind as Motion: Explorations in the Dynamics of Cognition* (Cambridge: MIT, 1995).