The Journal of Mind and Behavior
Winter and Spring 2012, Volume 33, Numbers 1 and 2
Pages 69–96
ISSN 0271–0137

Meditation on a Mousetrap: On Consciousness and Cognition, Evolution, and Time

Stephen E. Robbins

Fidelity Information Services

Evolutionary theory has yet to offer a detailed model of the complex transitions from a living system of one design to another of more advanced, or simply different, design. Hidden within the writings of evolution's expositors is an implicit appeal to AI-like processes operating within the "cosmic machine" that has hitherto been evolving the plethora of functional living systems we observe. In these writings, there is disturbingly little understanding of the deep problems involved, resting as they do in the very heart of AI. The end-state requirements for a system, device, or "machine" with intelligence capable of design are examined. The representational power must be sufficient to support analogical thought, an operation demanding transformations of events in imagery, in turn a function of perception, both dependent on a non-differentiable flow of time. The operational dynamics of the device must inherit this fundamental property of the dynamically transforming matter–field. Whether the evolutionary mechanisms or algorithmics thus far envisioned by biology or AI are coordinate with such requirements is left seriously in doubt.

Keywords: consciousness, artificial intelligence, evolution, time

Whether we are contemplating radios, robots, or robins, we are viewing very complex devices. For radios or robots, we know the device was created by human minds via a not well understood process called "design," and given the difficult birth of the radio, "creative design." For robins, the evolutionary theory of Darwin tells us things are different. The universe, acting as a giant machine, employed a form of procedure or "algorithm" to produce the robin. This procedure used random conjunctions of atoms to make chemical molecules. With more random conjunctions, it produced an elementary, living "device," perhaps a proto-cell. It then used and continues to use random mutations, in conjunction

This article is a modification of a chapter in Robbins (2012). Correspondence regarding this article should be addressed to Stephen E. Robbins, Ph.D., 2750 Church Road, Jackson, Wisconsin 53037. Email: serobbins27@wildblue.net

with forces or events in the external environment, to effect "natural selections" which dynamically transform devices into yet different devices, resulting in things such as robins, rabbits, and a Rex or two of the Tyrannosaurus type.

With this giant machine, we have removed all need to design these devices, and most significantly, any form of Mind or Intelligence designing them. This view is very much in consonance with Artificial Intelligence, which envisions machine algorithms that successfully design devices without any role required for consciousness, or conscious perception. The existence of AI and its mission is very much a hidden support of evolutionary theory. Indeed, Lloyd (2006) has proposed that the universe is a vast quantum computer wherein a few simple programs were constructed via random processes, enabling the bootstrapping of the whole complex production algorithm and machinery into existence.

There is a difficulty, however. Artificial Intelligence harbors a deep, unresolved problem, namely, that of *commonsense knowledge*. It is precisely this form of knowledge that underlies the construction of devices, be it mousetraps, mice, or mammoths. Knowledge of course is a function of mind. Mind, in turn, is an integral participant in a flow of universal time that is *indivisible* or *non-differentiable*. It is this simple fact that undermines AI's ability to solve the problem of commonsense knowledge, and as a result, any hidden support it could provide for the theory of evolution. In turn, this means that the Cosmic Evolutionary Machine must be a different "device" than that envisioned either by AI or by theorists of evolution.

It is not my purpose here to dispute the fact that there is evolution. But I intend to show that this extremely important subject, affecting profoundly our conceptions of man and mind, is being treated cavalierly by its expositors, and is far more complex than is being portrayed. In fact, we shall see that it is intimately entwined with this question: What is the relation of consciousness to cognition? Here, we shall see that our model of time is critical.

The Mousetrap and the Complexity of Devices

In recent years, consternation arose in the theoretical circles of evolution as Michael Behe (1996, 2007), an academic biologist, challenged the possibility of the "algorithmic" approach to design espoused by evolution. Though Behe dealt heavily in the biochemical realm, he placed the problem initially in the intuitive context of a mousetrap. The (standard) mousetrap consists of several parts (Figure 1). As a functioning whole, he argued, the trap is "irreducibly complex." For the device to work as designed, all the parts must be present and organized correctly, else it does not function.

The urge is to break the problem of instantiating this design into simpler components — evolving the separate, smaller parts. Natural selection buys nothing here, Behe argued. Natural selection picks some feature or form or

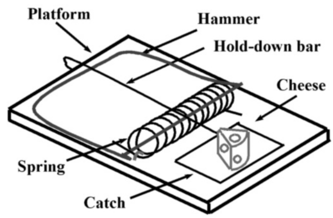

Figure 1: Mousetrap, standard issue.

component to continue because it happens to have been proven useful for survival. Evolving a single part (component), which by itself has no survival value, is impossible by definition — impossible, that is, by the definition of the role and function of natural selection. But even if by chance the parts evolved simultaneously, there remains the enormous problem of organization of the parts. How does this happen randomly? Each part must be oriented precisely spatially, fitted with the rest, fastened down in place, and even fabricated, etc. There are enormous degrees of freedom here — ways the parts can rotate, translate, and move around in space — which drive the odds against randomness to enormous proportions.

The problem can quickly be placed in the biochemical realm. Consider just one such structure in the cell alone. To manufacture palmitic acid, the cell relies on an elaborate circular molecular "machine." At the machine's center is a small arm comprised of molecules. The arm swings successively through six "workstations." Each time the arm rotates, two molecular subunits of the fatty acid are added by the action of enzymes at the workstations, and after seven rotations, the required fourteen units are present and the fatty acid released. For this rotary assembly to work, all six enzymes must be present in the right order and the molecular arm properly arranged. Now we ask, how, in what steps, always having a useful or survival value, does natural selection produce such a device?

Reviewers of Behe admit the lack of current solutions to this question. To quote one, "There are no detailed Darwinian accounts of the evolution of any fundamental biochemical or cellular system, only a variety of wishful speculations" (Shapiro, 1996, p. 63). Nevertheless, evolutionists have reacted strongly, with attacks focusing heavily on the biological and biochemical level. An interesting case is their attack upon a favorite example used by critics of evolutionary theory involving the gas-puff firing Bombardier beetle. The beetle (there are many variants) uses a chemical combination of hydroquinones and hydrogen peroxide which collect in a reservoir. The reservoir opens into a thick-walled reaction chamber (in the beetle's rear) lined with cells that secrete catalases and peroxidases. The resulting reaction quickly brings the mixture to a boiling point, vaporizing about a fifth. The pressure closes the valve and expels the gases through openings at the tip of the abdomen in a powerful jet at a would-be attacker. If the system were not initially designed with separate chambers for the chemicals, it is argued that the beetle itself would explode. The "exploding beetle" concept has been questioned, but more interestingly, Isaak (1997) has laid out a series of simpler beetle instantiations or steps, with examples of various steps embodied in other beetles of the class, which at least indicate a progression towards the Bombardier's sophisticated system.

In sum, there are definite biological arguments for the existence of simpler stages. Note, however, that while one can demonstrate that there are simpler stages, this does not mean that one has an actual, concrete model of how one transitions from stage A to stage B, and then to stage C. It was this that formed the implicit force of Behe's "irreducible complexity" argument (cf. Behe, 2007). At this point, evolutionary theory invokes natural selection, which chooses B over B' or B'', and which is effected by external forces of the environment. This is vague enough, while the actual creation of B, B', or B'' from A requires the mechanism of mutations.

That mutations can account for change in what is called "microevolution" is unquestioned. The fish in ponds in the depths of dark caves gradually turn white. Certain light-colored moths in England during the dusty, sooty era of the industrial revolution gradually turned to a darkish color. (With the decrease in industrial pollution, they have also recently "evolved" back again to a light color.) But the assumption has been that this same mechanism can work for larger, more complex, structural transitions, where we move from dinosaur to bird, fish to frog, frog to rat, or even from variant 1 to variant 2 to variant 3 of the Bombardier beetle. This is the point of contention, and here I must discuss things at the example level of the mousetrap.

The treatment of the mousetrap example per se by evolution theorists, with its question of transitions (from device A to device B, and from B to C), is less than satisfying. In fact, as we shall see, it actually moves in the realm of AI, a realm where there are great problems precisely in this design dimension. Keep in mind that while in the biological realm, we tend to talk about these transitions simply as "mutations," there is much more going on, for just as in the mousetrap, we are talking about complex spatial fittings and fastenings of parts, complex form shaping and fabrications of the parts from materials. To effect this, even considering the gene "switches" of "Evo Devo" (Carroll, 2005), would require extremely complex "programming" or modifications of the sequences in the genetic instructions to bring this about — i.e., long sequences of actions that must occur coherently, that leave random probability behind, and verge, at least, on *artificial* design.

Evolution Theorists Attack the Mousetrap

An argument, often cited as though it were a definitive critique, was provided by McDonald (2000) to demonstrate how the mousetrap could have simpler instantiations. His caveat is that this is not an analogy for evolution per se, but the argument *is* taken as a critique of Behe (e.g., Miller, 2003; Young and Edis, 2004). Working backwards, McDonald gradually simplified the trap, producing four "predecessor" traps of decreasing complexity. Behe argued, however, it is not that simpler mousetraps do not exist. The question is progression — the actual mechanism of movement from A to B to C. If McDonald *is* taken as a defense of evolution, Behe (2000) easily produces a strong counter argument. Starting with McDonald's first and least complex trap (Figure 2, left) in the "sort of evolving" series, he examined the steps needed for McDonald to arrive at the second trap (Figure 2, right). The first (or single piece) trap has one arm, under tension, propped up on the other arm. When jiggled, the arm is released and comes down, pinning the mouse's paw. It is a functional trap.

The second trap has a spring and a platform. One of the extended arms stands under tension at the very edge of the platform. If jiggled, it comes down, hopefully pinning some appendage of the mouse. To arrive at the second, functional trap, the following appears needed:

1. Bend the arm that has one bend through 90 degrees so the end is perpendicular to the axis of the spring and points toward the platform.
2. Bend the other arm through 180 degrees so the first segment is pointing opposite to its original direction.
3. Shorten one arm so its length is less than the distance from the top of the platform to the floor.
4. Introduce the platform with staples (neither existed in the previous trap). These have an extremely narrow tolerance in their positioning, for the spring arm must be on the precise edge of the platform, else the trap won't function.

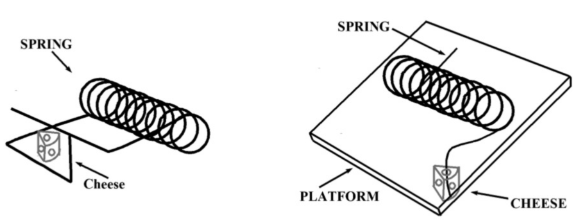

Figure 2: Mousetraps #1 (left) and #2 (right) from McDonald's (2000) first series. (All McDonald figures reprinted with permission.)

All of this must be accomplished before the second trap will function — an intermediate but non-functional (useless) stage cannot be "selected." This complicated transition is a sequence of steps that must occur coherently. With each step required, we decrease the probability of random occurrence exponentially.

Each of the subsequent transitions in the first series (2–3, 3–4, 4–5, where 5 is the standard trap) proved subject to the same argument. McDonald (2002) then produced a second, more refined series of traps. He argued that the point was made that a complicated device can be built up by adding or modifying one part at a time, each time improving the efficiency of the device. Yet there are still problematic transformations between many of his steps.[1] For example, in the second series, the transition between a simpler spring trap (Figure 3, trap five) and one now employing a hold-down bar (Figure 3, trap six) is a visual statement of the difficulty of the problem. Even if the simpler trap were to become a biologically based analog — a largish "mouse-catcher beetle" — sprouting six legs and a digestive system for the mice it catches, the environmental events and/or mutations which take it to the next step (as in trap six) would be a challenge to define.

But the most apparently decisive evolutionary argument is that indeed biological "parts" exist that in themselves are independently functional. In essence, then, evolution has available to it pools of independently functional components from which to select, and from which to build various larger functioning wholes. Kevin Miller (2003) considered this the finding of Melendez–Hevia,

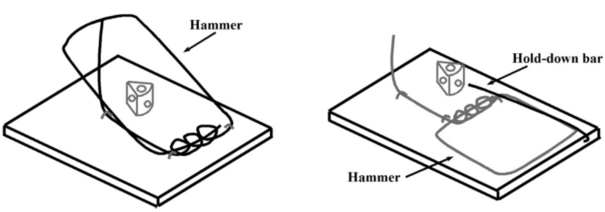

Figure 3: Traps five (left) and six (right) from the second series (McDonald, 2002). Trap six now has a hold-down bar hooked into the platform and lodged (lightly) under the hammer arm.

---

[1]Because (for example) simpler mousetraps are shown to exist, irreducible complexity is critiqued as vague. The two traps of Figure 3, however, clarify the issue. Trap five is simpler than trap six. But each trap is irreducibly complex; each fails to work as designed without all its components. In some cases, the trap is indeed a slightly simpler version of the same design, as trap six of Figure 3 might be taken as a simpler version of a standard mousetrap which works without one of the standard trap's parts. But inevitably the simpler traps morph to different designs which no longer effect quite the same function (e.g., trapping a paw vs. smashing the poor creature).

Waddell, and Cascante (1996) in the realm of the Krebs cycle.[2] Miller applies this logic to the mousetrap. Each component can be conceived to be an independently functional part. For example, the hold-down bar can serve as a "toothpick," the platform as "kindling," three of the components can work together as a "tie clip" (platform, spring, and hammer), and so on. The implication of this argument is disturbing, for it indicates that the grasp of the problem is deeply insufficient. Either the evolutionists, at this point, have simply become very weak AI theorists, or they know something the AI folks don't know. The fact is, evolution theorists have blundered into the greatest of unsolved problems in AI, that of *commonsense knowledge*.

The Problem of the Mousetrap

Ironically, my own intellectual career had an early phase wherein I contemplated what it would take for an AI program to design a mousetrap (Robbins, 1976). The problem was presented as an initial list of components. For example, and not exhaustively, a 12" cubical box, a sharpened pencil, a razorblade, a length of string, paper clips, rubber bands, staples, toothpicks, and of course a piece of (Wisconsin) cheese. From this, the task is to create a mousetrap. (At the time, I believe, this was used as a creativity test for future engineers.) One AI program I considered was Freeman and Newell's (1971). This program had a list of *functional requirements* and *functional provisions* for various objects. For example, to design a KNIFE, it discovered that a BLADE *provided* cutting, but *required* holding. A HANDLE *provided* holding. By matching the requirements to an object's list of provisions, the program "designed" a knife. It is precisely the implicit approach of Miller (2003), as noted above.

I tried mightily to imagine how such a program would work in the mousetrap problem. There are many possible designs. I might make a form of crossbow, where the ends of the rubber band are attached to the outside of the box, the pencil (as an arrow) drawn back through a hole in the side, a paperclip holds it via a notch in the pencil, and a trip mechanism is set up with the paperclip, the string, and cheese. Or I might devise a sort of "beheader," where the razorblade is embedded in the pencil as an axe, the pointed pencil end lodged in a corner, the whole "axe" propped up by a toothpick with downward tension from the rubber band, string attached to the toothpick for a trip mechanism, etc.

What, I asked, would the database of objects' functional provisions and requirements look like? To make the story short, I will say that I quickly abandoned any hope for this scheme. The problem is far larger. One rapidly starts

[2]Behe, however, notes that this is simply like describing the various chemical transitions of oil, from its initial raw state, to gasoline, while ignoring the origin and explanation of the various and complex machinery employed at each stage of the refinery process.

to entertain the storage of "features." Noticing the "sharpness" of the pencil, it seemed, was integral to seeing it as supportive of the killing-function within the crossbow. It is doubtful that "killing" or "piercing" would have been listed in the database as "functional provisions" of a pencil. The corner of the box provided "holding" for the pencil-axe, and while it is doubtful this would have been listed as a functional provision of box corners, it seems a type of feature. Note, meanwhile, that in the axe case, the pencil "provides" something quite different from the pencil as arrow, while a certain feature of strength and rigidity has emerged in this context.

So do we envision a vector of pre-defined "features" for each object in our database? At a later date, in essence, this would be the approach of Gentner (1983) and many subsequent connectionist instantiations (Doumas, Hummel, and Sandhofer, 2008; Holyoak and Thagard, 1997; Hummel and Holyoak, 2005). But features are very ephemeral — they are functions of *transformations*. A fishing rod can be flexible under one transformation, sufficiently rigid under another. A floppy sock, under the appropriate transformation, gains sufficient rigidity to become a handy fly-swatter. The pencil's rigidity under one transformation may change to just enough flexibility to support the launching of spit wads. A box may preserve its edges and corners invariant under various rotations, but lose them completely under a smashing transformation applied by the foot. And precisely the latter may be done to turn the small box in the potential components list above into a temporary dustpan. Thus we would need to store all possible transformations upon any object.

Transformations

McDonald (2000), as we saw, performed two "bending" transformations on the wire of mousetrap #1 to obtain mousetrap #2. This form of dynamic transformation in thought heavily impressed the Gestalt psychologist, Max Wertheimer (1945). He had observed children in a classroom being taught, via drawings of a parallelogram on the blackboard, the traditional, algorithmic method of dropping perpendiculars to find the area. Yet, when Wertheimer himself went to the board and drew a rotated version of the parallelogram figure, he was shocked to see that the children failed to extend the method. But outside the algorithmic-oriented classroom, Wertheimer observed a five year-old who looked at a cardboard cutout of a parallelogram, then asked for a scissors so she could cut the (triangular) end off and move it to the other side to make a rectangle. This was bettered by the dynamic transformation exhibited by another five year-old child who folded the cardboard parallelogram into a cylinder, then asked for a scissors to cut it in half, announcing it would now make a rectangle.

We meet this dynamic "folding" transformation in Penrose (1994). While his critique of AI was heavily attacked by the AI community, few noticed that in his characterization of "non-computational" thought, Penrose had gravitated towards transformations and the invariants preserved under these transformations. In his proof that successive sums of hexagonal numbers are always a cubical number (hence a computation that does not stop), he initially folds a hexagonal structure into a three-sided cube. He then has us imagine building up any cube by successively stacking (another transformation) these three-faced arrangements, giving each time an ever larger cube (Figure 4). This is a dynamic transformation over time, in fact multiple transformations with invariants across each. We can expand the hexagonal structures successively, from 1, to 7, to 19, etc., each time preserving the visual hexagonal invariant. Then, each is folded successively, each time preserving the three-faced structural invariant. Then imagine them successively stacking, one upon the other, each operation preserving the cubical invariance. Over this event, the features (or transformational invariance) of the transformation are defined.

These cases are images of *events*. It is the ability to represent events in the medium of an image that has been so problematic to the information systems approach in cognitive science. Pylyshyn (1973) initially denied any need for mental images, arguing that the information in data structures is entirely sufficient to subsume the function of images. Later, in his "null hypothesis," while not denying their existence, he challenged the field to explain why images are needed. His key question was this: "What does the real work in solving the problem by [mental] simulation — a special property of images . . . or tacit knowledge?" (Pylyshyn, 2002, p. 162). Thus, in contemplating the folding experiments of Shepard and Feng (1972), where subjects were required to mentally fold paper into objects of certain forms, he noted that the subjects had, by necessity, to proceed sequentially through a series of folds to attain the result.

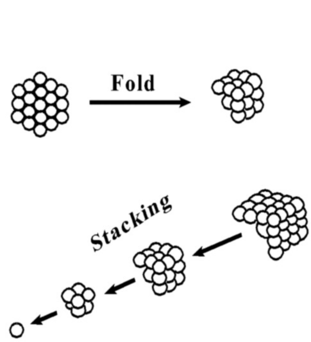

Figure 4: Top: A hexagonal number (19) form folded into a three-faced (side/wall/ceiling) structure. Bottom: Successive cubes built from side, wall, and ceiling. Each side, wall, and ceiling structure make a hexagonal number.

Why? "Because," he argued, "*we know what happens when we make a fold*" (2002, p. 164, original emphasis). It has to do, he stated, with "how one's knowledge of the effects of folding is organized" (p. 164).

Sloman (1971), in a seminal paper, had already given Pylyshyn his answer. He contrasted the Fregean or syntactic mode of representation with what he termed the analogic mode. In the analogic mode, there is the natural representation of *constraints*. The paper does not disintegrate while it is being folded. The edges stay stable and move to overlap one another. One surface generally stays stationary. All these constraints are in fact invariance laws defined over these event-transformations. On the other hand, in syntactic systems, failures of reference are commonplace. The syntactically correct, "The paper screeched and burbled as it was folded," makes little semantic sense — it instantly violates the invariance across folding events. The frame problem (McCarthy and Hayes, 1969) is in essence another statement of this problem of representational power (Robbins, 2002). To Sloman, the greatest challenge faced by AI was achieving this (analogic) form of representation.

Again, we can recast Sloman's challenge: What type of "device" is required to support this form of representational power? But this is only to ask: What type of device can support perception? No visual imagery ever occurred without visual perception. The congenitally blind bear witness to this. The image is a question of (1) perception and (2) the memory of this perception. In turn, the image *is* the knowledge. It is no less the knowledge than the actual perceiving of an event of folding is simultaneously — knowledge. What is a "fold" other than an invariant defined over transformations in concrete experience? We have seen folds made in sheets, folds made in paper, folds made in arms/elbows, folds made in sails, folds made by Penrose (1994) in three-faced hexagonal structures to make partial cubes, and even folds made with poker hands. And we have made the folds with bodily action. *Something* is always being folded. There is no such thing as an abstract "folding," no such thing other than as a dynamic transformation preserving an invariant and defined over our concrete, perceptual experience.

The Invariance Structure of Events

Transformations and invariance — why the emphasis? Firstly, discovering invariance laws *is* scientific explanation. This has been heavily argued (Hanson, 1958; Kugler and Turvey, 1987; Wertheimer, 1945; Wigner, 1970; Woit, 2006; Woodward, 2000, 2001, 2003). In this, science only models itself after the brain in perception. $E = mc^2$ is an invariance law. $F = -kX$ is an invariance law. In relativistic physics, it is only the invariants ($d = vt$, $d' = vt'$) that are the realities of the relativistic universe (Lieber and Lieber, 1945), for it is these that hold across space–time partitions. This essential endeavor of science is often beclouded in the psychological sciences, but it is invariance laws that characterize the ever transforming world of perception where events occur in the concrete ecological world. As I have stressed many times (Robbins, 2002, 2004a, 2006a, 2006b, 2007, 2008, 2009), such events have an *invariance structure*. An invariance structure is defined as such: *the transformations and invariants specifying an event and rendering it a virtual action*.

A simple event that is illustrative is stirring coffee. The swirling coffee surface is a flow field (Figure 5), in this case in radial form. The constant size of the cup, as one's head moves forward or backward, is specified, over time, by a constant ratio of height to the occluded texture units of the table surface gradient. Over this flow field and its velocity vectors a value, τ, is defined by taking the ratio of the surface (or angular projection) of the field at the retina, r(t), to its velocity of expansion at the retina, v(t), and its time derivative. This invariant, τ (or tau), specifies time to impending contact with an object or surface, and has a critical role in controlling action (Kim, Turvey, and Carello, 1993). A bird, for example, coming in for a landing, must use this τ value to slow down appropriately to land softly. As the coffee cup is moved over the table towards us, this value specifies time to contact and provides information for modulating the hand to grasp the cup (Savelsbergh, Whiting, and Bootsma, 1991). As the cup is cubical, its edges and vertices are sharp discontinuities in the velocity flows of its sides as the eyes saccade, where these flows specify, *over time*, the form of the cup (Robbins, 2004a, 2007). The periodic motion of the spoon is a haptic flow field that carries what in physics is termed an *adiabatic* invariance — a constant ratio of energy of oscillation to frequency of oscillation (Kugler and Turvey, 1987). The action of wielding the spoon is defined by an inertial tensor, the diagonal elements of which represent the forces involved, or more

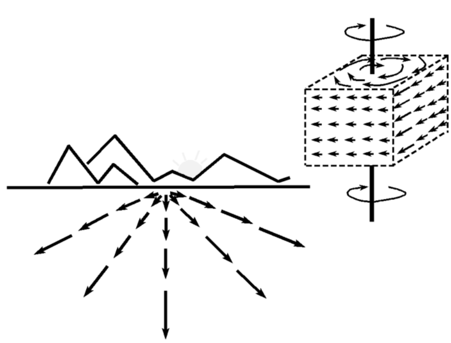

Figure 5: Optical flow field. A gradient of velocity vectors is created as an observer moves towards the mountains. The flow field "expands" as the observer moves. At right, the flows as a cube rotates towards the observer.

precisely, the object's resistance to angular acceleration (Turvey and Carello, 1995). This entire structure and far more must be supported, globally, over time, by the resonant feedback among visual, motor, auditory, even prefrontal areas. In other words, *it is this entire informational structure that must be supported, in ongoing fashion, over time, by the neural dynamics supporting the perception of the coffee stirring event.* It is in these invariance structures that we find the foundation of knowledge and semantics (Robbins, 2002, 2008). Knowledge and semantics are both served by a fundamental memory operation termed *redintegration*.

Redintegration, Commonsense Knowledge, and the Frame Problem

As I am walking along a road, I spot a rustle in the grass in the roadside embankment. Instantly an experience returns in which several blacksnakes rushed by me as I was walking up a hill years ago. This is the elementary operation of redintegration. It is the most ecological of memory operations. Wolff (1732/2010), a disciple of Leibniz, first coined this law in 1732 in his *Psychologia Empirica*, stating that "when a present perception forms a part of a past perception, the whole past perception tends to reinstate itself." Klein (1970) notes that these remembered experiences are "structured or organized events or clusters of patterned, integrated impressions," and that Wolff had in effect noted that subsequent to the establishment of such patterns, the pattern might be recalled by reinstatement of a constituent part of the original pattern. It is the mathematical description of these "event patterns" in terms of invariance laws that is the core of Gibson's theory.

The redintegration principle can be stated simply:

> An event E' will reconstruct a previous event E when E' is defined by the same invariance structure or by a sufficient subset of the same invariance structure.

I will not discuss in this paper how a time-extended experience is "stored." It is sufficient here to assume the principle of exemplar theory (Crowder, 1993; Goldinger, 1998), which holds that every experienced event, in every detail, is stored. Given the discussion above, this means the event's entire time-extended dynamic structure with defining invariants. While exemplar theory simply uses the vague notion of events leaving "traces," we can simply envision, as did Gelernter (1994), a "stack" of experienced, coffee stirring events in memory, in fact, every coffee stirring event ever experienced. When a present event, E', is perceived, with the brain therefore supporting the time-extended invariance structure of E', we can say with exemplar theory, that all the previous event traces are activated, or, more accurately, that this entire stack of experiences with similar dynamic structure, is resonant with E'.

We can imagine, then, a robot stirring coffee. As he stirs, the coffee liquid medium begins to behave as a thick cement, barely allowing the spoon to be in motion. This is just one of a vast list of possible anomalies. For others: as the robot stirs, the cup floats off the table, or the motion of the liquid is in a counter-circular direction to the spoon, or the cup bulges in and out, or small geysers erupt from the liquid surface, or the sound is a "snap, crackle, pop" like Rice Krispies, or the spoon melts into rubber . . . . While it is not uncommon to see philosophers discussing the problem in terms of the robot updating his "beliefs" about coffee stirring, this is misleading. The prior, far more fundamental question is this: How does the robot detect that this (or any of the above) is an unexpected feature of the event? In the context of the frame problem, as the event is ongoing, the robot must check, continually, his vast list of frame axioms defining not only the features of this event, but multitudinous dimensions of his external world. Discovering a method to reduce the list of axioms is exactly the frame problem.

In redintegration, we obtain a view of a far more powerful method. The anomalous stirring event, with cup bulging in and out, retains sufficient invariance structure to send a redintegrative cue throughout the "stack" of stirring experiences, retrieving similar events of coffee stirring. Yet there will be an "interference," a dissonance with the whole. Since we are dealing with a very concrete "device," it is a *felt* dissonance — the discrepancy is instantly detected — and there is no need to check a list of frame axioms to see if this is an unexpected feature of the event.

As the body/brain is such a redintegrative device, this is, in essence, its method of solving the frame problem. Within this method, there lies implicitly its approach to the correlated problem of commonsense knowledge, and therefore the design of devices. I will develop this in what follows and as we examine the approach of AI and cognitive science to this problem.

Connectionism versus Ecological Invariance

Connectionist models propose to be presenting the method by which the brain represents semantic knowledge or semantic cognition. Rogers and McClelland (2004, 2008) present a scheme using a three-layer network. The input units correspond to an item in the environment, for example: ROBIN, or SALMON, or FLOWER. The units in the relationship layer correspond to contextual constraints on the kind of information to be retrieved, for example: IS, CAN, HAS. The input pair, ROBIN CAN, they argue, corresponds to a situation in which the network is shown a picture of a robin and asked what it can do. The network is trained to turn on the correct attribute units of the output layer, in this case: GROW, MOVE, FLY, SING (as opposed to SWIM, DIVE, FLOP). As the connection weights are initially random, the output units of the network must be adjusted gradually, via a backpropagation algorithm based upon the amount of error relative to the desired output state. This adjustment often requires many hundreds of epochs of training.

Rogers and McClelland hold that this network is perfectly at home in the ecological world. The input units, they hold, can be construed as receiving perceptual input, for example the observation of a robin sitting on a branch, and the output units are predicting possible events or outcomes, say, the robin flying away. Obviously these statements would hold for, "The rustle in the grass" that predicts (retrieves) the slithering snakes. As is the norm in these models, no effort is made to determine if this network could actually support the complex patterns that we have seen characterize time-extended events, or also, problematically, whether it makes any realistic, ecological or evolutionary sense to demand of the model of the brain that supports this form of redintegration that it require hundreds of epochs of training to establish this memory relationship.

We can place another ecological learning situation within the Rogers and McClelland framework. Given the object, SPOON, in the context of CAN, the network would be trained to respond with the set of things a spoon can do, for example, STIR (as in coffee), SCOOP (as in cereal), CUT (as in grapefruit), BALANCE (as on the edge of the coffee cup). In essence, for an event such as stirring coffee, we have bifurcated these various events into components — SPOON and STIR, or SPOON and CUT, or SPOON and SCOOP — and attempted to train the network to *associate* these components.

What sense does this make? In reality, we are perceiving the spoon as an integral part of a stirring event, with all the event's ongoing invariance structure, to include the forces supplied by the spoon relative to the liquid medium, the resistance of the medium, its particular motion, the periodic motion of the spoon with its inertial tensor, adiabatic invariance, etc. It is a structure that is necessarily being supported, over time, by the neurodynamics of the brain, else there is no perception of the ongoing event with its structure (Robbins, 2008). Where is the "error?" That is, where is the error that must be weight-adjusted to achieve the proper "linkage"?

The fact is, this partitioning of events into arbitrary components harkens back to Ebbinghaus, who made the move of removing all semantics from the study of memory, inventing instead, the nonsense syllable. When studying how we learn nonsense syllable pairs such QEZ–WUJ, memory research is being faithful to this vision — studying the process of the formation of the elementary item-bond. The unceasing desire to explain this "bond" is the elementary ill of associationism. Subjects in these experiments quickly learned that if they could form an event, say a pudgy (pudge for WUJ) Turkish person wearing a fez (fez for QEZ), they could learn the pairs more easily. Paivio's (1971) introduction of imagery into these experiments was the first near-ecological crack in the approach. For an arbitrary pair such as DOG–GATE, the subjects now imagined an event such as a dog opening a gate, and performance greatly improved. The connectionist net is learning syntax-rules. Syntax can be defined as *rules for the concatenation and juxtaposition of objects* (Ingerman, 1966). QEZ–WUJ is a rule for the juxtaposition of objects, as is DOG–GATE when DOG and GATE are treated at the merely mechanical level as a pair of "marks" or objects, as is SPOON–CAN–STIR, etc. The redintegrative process described above relies, rather, on the laws (invariance structure) of events.

But suppose we have error-trained the connectionist net such that for SPOON CAN, it responds with the set: STIR, SCOOP, CUT, BALANCE. This represents the network's semantic "understanding" of the capabilities of a spoon. But we can easily understand the sentence: The SPOON CAN CATAPULT (a pea). We understand this because we grasp that the spoon will support the forces/invariance structure of catapulting. It is in the invariance structure that the semantics of this sentence rests.

The difficulty for the connectionist net rests precisely in the realm of the powerful critique made by French (1990) in the context of the Turing test. French proposed various tests for any computer attempting to masquerade as a human. Obtaining a passing grade relied totally on having the requisite concrete experience. One test was a rating game, with questions such as:

> Rate purses as weapons.
> Rate jackets as blankets.
> Rate socks as flyswatters.

And of course we could have:

> Rate spoons as catapults.

The computer's ratings would be compared to human rating norms. French argued that there is no way a computer can pass such a test without the requisite concrete experience. The problem equally holds for evaluations of statements such as:

> A credit card is like a key.
> A credit card is like a fan.

The list is endless. Says French (1999), ". . . no a priori property list for 'credit card,' short of all of our life experience could accommodate all possible utterances of the form, 'A credit card is like X'" (p. 159). Without the experience, one incurs the necessity of either pre-programming or training-up the association weights of all possible pairs of objects. Yet, this is exactly the implicit road down which the network of Rogers and McClelland is headed. To even bring

SPOON into some form of association with CATAPULT would require additional, explicit epochs of weight adjustments involving CATAPULT. But a catapult is just one of a vast array of objects we could "associate" with a spoon. We could, for example:

Rate a knife as a spoon.

A knife can serve as very good stirrer of coffee, showing the structural invariance required to move the medium under this motion — if this is the transformational context. It is not much good for eating soup. But this makes the programming of association weights even more impossible, for now they all depend upon a transformational context. As French essentially noted, the neural net has no concrete experience with stirring, spoons, knives, or catapults. But what is experience? At minimum, it is comprised of multimodal events structured by time-extended transformations and the invariants preserved over these.

The rating events above are all forms of *analogy*. In each, we have, in effect, *the projection of an invariance structure upon a possible component*. A knife is placed in a stirring event, a spoon in a catapulting event, a sock in a fly swatting event, or a box and pencil in a beheading event, and each "tested" on the emergence of the structural invariance ("features") requisite for preserving the invariance structure of the event. Here, *the analogy defines the features*.

AI's Approach to Analogy

The symbolic programming method in AI has proffered several models for analogy making, the most famous of these being Gentner's (1983) Structure Mapping Engine. To the Structure Mapping Engine, as in all AI, *the features define the analogy*. Thus the Structure Mapping Engine treats analogy as a mapping of structural relations relative to pre-defined features. The solar system, for example, and the Rutherford atom both have specific features and their relationships described in predicate calculus form, e.g., Attracts (sun, planet), Attracts (nucleus, electron), Mass (sun), Charge (nucleus), etc. Chalmers, French, and Hofstadter (1992) level a heavy critique upon this approach, noting the helplessness of the Structure Mapping Engine without this precise setup of features and relations beforehand, and with this setup given, the purely syntactical, nearly "can't miss" algorithmic or "proof" procedure that follows. The resultant discovery of analogy is, to quote these critics, a "hollow victory."

The connectionist models of analogy are equally wedded to this approach. For Discovery of Relations by Analogy or DORA (Doumas, Hummel, and Sandhofer, 2008), the engine for forming analogical relations is a comparator that operates on propositions which have a dimensional value. If DORA "thinks" about a DOG of size-6 and a CAT of size-4, the comparator, detecting the dimensional value, links a "more" relation to the size-6 (or "*more*+size-6") related to the DOG and a "*less*+size-4" for the CAT. If this pattern reminds DORA of a previous comparison of the same type between a BEAR (more+size-9) and a FOX (less+size-5), a further operation now compares the CAT and DOG units to the similar setup for the BEAR and FOX, eventually spawning a new unit, BIGGER, bound to BEAR and FOX or *Bigger* (BEAR, FOX). The authors of DORA argue that this same process will be fully applicable to ecological events, i.e., "relations" such as "chasing," and by extension, "stirring."

Ignoring for the moment that DORA's comparator is not even close to something that can handle actual, ecological events, let us suppose we have formed single place predicates (SPs) such as *stirred* (coffee), *stirrer* (spoon), and *stirred* (paint), *stirrer* (paint-stick). According to the model, a pair of single place predicates enters working memory, in this case *stirred* (coffee) and *stirrer* (spoon). These are "mapped" as a unit onto other SPs, in this case *stirred* (paint) and *stirrer* (paint-stick). This mapping serves as a signal to link the SPs into a larger predicate structure, thus *stir* (spoon, coffee) and/or *stir* (paint-stick, paint).

This is simply a syntactic mapping. It is based on the fact that the model would attach "stirrer" as a feature to spoon, and "stirred" as a feature to coffee. Given the precise setup of these predicates, the mapping can occur via an algorithm. Without this precise setup, the process is helpless. The network has no ability to create or recognize the validity of multi-place predicates such as *stir* (knife, coffee) or *catapult* (spoon, pea) without this setup. It is another example of the validity of French's critique. There is nothing in the network, unless it has been specifically trained and the "features" specifically set up, that would support these relations. Connectionism has simply met symbolic AI at the same problem — commonsense knowledge.

At DORA's heart is a model of redintegration. When DORA envisions *stirred* (coffee) and *stirrer* (spoon) entering working memory while other propositions in long term memory that share the semantic units — *stirred* (paint) and *stirrer* (paint-stick) — are brought in and made available for mapping, this is the redintegration of events. DORA's is based on a very problematic reactivation of the "same" semantic units. Underlying the stirring of paint-sticks, spoons, or spatulas are the complex invariance laws we have seen described — the "wielding" characterized by inertial tensors, the adiabatic invariance underlying the periodic motion, the radial flow field of the liquid's surface. There are no simple dimensional values analogous to "size-6," for example a "wielding-6" or "periodicity-3," that can be assigned to "semantic units" such that these that can now be "compared" via the simple algorithm of DORA. DORA has no ability to deliver on what "same" can possibly mean in these kinds of ecological events, for DORA, as in all connectionist approaches, utterly begs the description of change.

The invariance structure of the event is the description of change. It is this underlying structure that would need to be invoked as a constraint to prevent

DORA from "thinking about" stirrer (spoon) and chased (Mary), and being reminded of a previous comparison, stirrer (spatula) and chased (Joe), thus deriving stirring (spatula, Joe). The proposition, stirring (spatula, Joe), is the essence of a syntactic "failure of reference." As a sentence, it takes its place with other sentences that are syntactically correct but seem to have no semantic justification:

1. The leaf attacked the building.
2. The shadows are waterproof.
3. The spatula stirred Joe.
4. The building smoked the leaf.

Katz and Fodor (1963), early in the game, tried to solve this problem by "semantic markers" assigned to each lexical item in the deep structure. These were simply syntactic rules trying to represent physical constraints — rules attempting to do the work of the invariance structure. The "leaf " in (1) would thus receive a marker denoting it as *inanimate* among other things, while "attack" would receive a marker requiring its use with an *animate* object. Having incompatible semantic markers, such a system brands the sentence as meaningless. "Stirring" would have been tagged with a marker requiring its object to be, say, liquid. Joe, having no such marker, would have thus been seen as illegal in (3) and the string also branded as meaningless. Unfortunately such sentences can appear very meaningful. An analogy performs a transformation; it allows the requisite "features" to emerge. Sentence (2), which would also have incompatible markers, is perfectly interpreted as meaning that we can throw as much water on shadows as we like and they will be unharmed, i.e., the perceived event of water pouring upon a shadow shares an invariant with other events of water pouring over waterproof materials, namely the undamaged state of the material substance of these objects under this transformation. As for (3), we *can* easily make sense of this sentence, "The bad architecture of the system is like a spatula, stirring Joe, the programmer, into an anxious mess." Such transformations would quickly lead to "rules for relaxing the rules," but the rule system quickly ends in anarchy, being so flexible that it is useless as an explanatory device.

The apparent meaninglessness could only be avoided by a constraint, but this constraint is equivalent to having — stored somewhere and acting — the complete invariance structure of the event (of stirring, of pouring, etc.)! The invariance structure is what prevents Joe from "being stirred" given the normal context of a stirring event — Joe is not easily inserted into this dynamic structure. It is this structure that causes the feeling of anomaly — the failure to resonate with the laws of experience — in the sentence:

As Joe stirred, the coffee snapped, crackled, and popped.

This structure cannot be syntactically represented. You are begging an entirely different form of knowledge to supply the vast number of possible constraints involved even in this simple event. It is the event invariance structures that are prior in explaining these linguistic cases.

The features on which analogy is "based" cannot be preset, pre-defined. As noted, it is the analogy that defines the features. Analogy is a transformation. This is to say that it is a process that occurs over a concrete flow of time. It is supported only over concrete experience or the remembrance thereof, i.e., it is carried only over transforming imagery — the figural mode. Artificial intelligence, based in a classical notion of an abstract, spatialized time and without a theory of perception, can support neither of these requirements for analogy, and it is analogical thought that is supporting the design of the mousetrap.

Beyond the Fundamental Metaphysic of AI

AI is founded in what can be termed the classical metaphysic. It is the same metaphysic that lurks beneath the hard problem/the origin of qualia. I have laid out arguments several times (Robbins, 2000, 2001, 2002, 2004a, 2004b, 2006a, 2006b, 2007, 2008, 2010a, 2012) on the consequences of this framework and on the alternative model that exists in Bergson (1896/1912) when combined with Gibson (1966). The essence of this classic metaphysic is an *abstract* space and time. The space is conceived as continuum of points or positions. Time is simply another dimension of this space. Thus the motion of an object (itself a set of points) in this continuum is treated as a movement from (static) point to (static) point along a line or trajectory. This is an infinite regress, for to account for the motion, we must reintroduce yet another line/trajectory of points between any two adjacent static points on the original line, ad infinitum. This spatial treatment of motion is the origin of Zeno's paradoxes — the arrow, always occupying a static point in the continuum, "that never moves," or Achilles, forever halving/dividing the distance, who never catches the hare. Indeed, for Bergson, this space is simply "a principle of infinite divisibility."

Bergson argued that to escape this, we must treat motion as indivisible, or as Nottale (1996) now states it, as *non-differentiable*. Motion is better conceived as a melody where each note ("instant") interpenetrates the next, and each is the reflection of the entire preceding series — an organic continuity. As the object can move across the continuum, or the continuum (or the coordinate system) can be moved beneath the object, all motion becomes relative; all *real* motion is now lost. But stars die, trees grow, couch potatoes get fat — there must be real motion. Rather than "objects" in motion, we now view the *whole* of the matter–field as transforming, where the motions of "objects" are now *changes or transferences of state.*

As opposed to the (quality-less) homogeneity inherent in the abstract continuum of mathematical points, the matter–field is now intrinsically qualitative, and the nature of its non-differentiable motion gives the entire universal field, in its time-evolution, a fundamental property of memory. Each “past” instant does not recede into non-existence as the “present” instant arrives. This “primary” memory inherent in the indivisible motion of the field makes possible the brain’s specification of a past history of the motion of this qualitative field — a rotating cube, a buzzing fly, a folding hexagon, or a bending mousetrap arm.

In the context of this “specification,” I have given arguments for ceasing to view the world as being encoded or represented within the brain, and seeing the brain, rather, as itself the decoder. The decoding is effected by the brain in the role of a concrete reconstructive wave passing through the external, holographic matter–field, with the brain’s state being specific to a past motion of the field. Via the brain’s energy state (or its underlying chemical velocities), it is a specification at a particular scale of time or in essence a space–time partition — a “buzzing” fly as opposed to a fly flapping his wings like a heron. The “image” (of the fly) is not mysteriously generated by the brain; it is now simply a *diminution* of the whole, a specification of a subset of the vast information in the dynamically changing holographic field. The brain is not simply a “hologram.” The reentrant neural processes, the oscillations, the resonant feedback that have hitherto been taken solely to be abstract computations — all in effect contribute to this very concrete wave. The brain’s function is as concrete as that of an AC motor. The motor creates an electric field of force; the brain creates a concrete, continuously modulated reconstructive wave “passing through” the matter–field.

The modulation pattern is driven by the invariance structure of the external events in the ecological world. It is the invariance laws defining events that drive what the brain, as a reconstructive wave, specifies as the external image. As in relativity, we require invariance laws, for it is such laws that hold across possible scales of time or space–time partitions. The specification is always an *optimal* specification based on the probabilistic information — with its inherent uncertainty due to the continuous flux of time (Lynds, 2003) — available to the brain. Even illusions are optimal specifications of a past form of motion of the matter–field.

Five Requirements for an Embedded Intelligence

In this context, we can derive five requirements for a device that supports perception, and therefore cognition, and thus, the ability to design:

1. The total dynamics of the system must be proportionally related to the events of the matter–field such that a scale of time is defined upon this field.

2. The dynamics of the system must be structurally related to the events of the matter–field, i.e., reflective of the invariance laws defined over the time-extended events of the field.
3. The operative dynamics of the system must be an integral part of the indivisible, non-differentiable motion of the matter–field in which it is embedded.
4. The information resonant over the dynamical structure (or state) must integrally include relation to or feedback from systems for the preparation of action (for from the vast information in the holographic field, the principle of selection is via relation to possible action by the body).
5. The global dynamics must support a reconstructive wave.

To support perception, then, the device (and its "processing") must literally be embedded in the non-differentiable time-flow of the matter–field. A syntax-directed processor does not meet this requirement. Though it is felt by some (Dietrich and Markman, 2000; Prinz and Barsalou, 2000) that the operations of a computer riding atop its continuous dynamics can support semantics (and by implication experience and perception), this is not the case, and it is why, in (3), the term "operative dynamics" is used. In the computer model, the effective, operative "dynamics," if you can call it that, is in the syntactic manipulation of symbols. The concatenation and juxtaposition of objects in the classical abstract space and discrete-instant "time" — operations, further, for which the scale of time is utterly irrelevant — is not sufficient to support perception or the continuous, time-extended transformations characteristic of analogical thought.

And in general, it is not just the organization of components, or the material from which they are made. It is the concrete *dynamics* they support. As Haselager (2005) notes in the context of supporting an autopoietic system, "You cannot make a boat out of sand." Neither does one create the concrete, electric wave of an AC generator with the "proper organization" of toothpicks, rubber bands, or abacus beads. Whether biological or artificial, the dynamics required for perception must support a very concrete wave, establishing a ratio of proportion, i.e., a scale of time, upon the matter–field. It is this fundamental architecture that is required to support the time-extended images of perception, and therefore the time-extended, transforming images of memory employed in analogical thought.

The Broadly Computational Mousetrap

We return then to the "device" underlying design. In the mousetrap task, we are designing from existing materials. I do not say from "existing components" because none of the objects is yet true a component, though each has an independent function (e.g., a pencil, a rubber band). The invariance structure of an event — the drawing back and firing of a crossbow, the striking down of the axe — is being projected over the possible "components." In the process, their requisite features emerge.

This is a powerful transformation over a non-differentiable time. I have striven here and elsewhere (Robbins, 2002, 2006a, 2006b, 2012) to lay out the basis for a device with sufficient representational power to support it and the implications for cognition it contains, to include the origin of the compositionality and systematicity required by Fodor and Pylyshyn (1988), the origin of the symbolic, and the nature of *explicit* memory and thought (Robbins, 2009, 2012). As Penrose argued, it is not computational in the abstract sense given by Turing. Turing's definition is predicated upon the abstract space of the classic metaphysic; it captured the mechanical computations of the bank clerks of Turing's 1940s era, or the mechanical knowledge and calculations of the parallelogram-challenged children in Wertheimer's classroom (Robbins, 2002). It did not capture the computation of the five-year old who dynamically transformed the cardboard parallelogram into a cylinder. The manipulation of discrete symbols in an abstract space and time cannot support this, nor will a dynamical device that cannot support perception. Rather, the dynamical brain or robotic system must generate a very concrete waveform in concrete, non-differentiable time, a wave which supports a broader form of computation, broader than Turing's narrow definition, but consonant with a broader definition he left fully open (cf. Copeland, 2000; Robbins, 2002).

Evolutionary AI

I am led to the conclusion that a "device" of this power, inheriting attributes of the non-differentiable time-flow of the matter–field in which it is embedded, is required to support the design transitions posed by McDonald's mousetraps. AI, in its current form, is far from the basic requirements for an intelligent device described above. Evolution theory cannot implicitly rely on AI-like algorithms for producing forms and creatures, whether mousetraps, mice, or beetles; it cannot rely on Lloyd's (2006) giant, cosmic quantum computer — a computer, no matter how quantum, that is still in the Turing class of computing machines.

Now, of course, evolutionary theory says that it does not rely on AI. It puts its weight on natural selection and mutations. To be clear, it must put *all* its weight on natural selection together with mutations (or "variation"). I am simply removing any temptation to go beyond this. Unfortunately, evolution's expositors have already succumbed to the temptation. Not even counting Lloyd's explicit appeal to programs underlying evolution, here is an example: bacteria have a "flagellum" — a thread-like propeller that drives them though the water. This little device has a rotating axle, turning inside a bearing, driven by a molecular motor. Behe thought it another irreducibly complex device. Dawkins (2006), while ridiculing Behe to the point of impugning his motives for publishing, approvingly references Kevin Miller — the same Kevin Miller who saw no problem building mousetraps from arbitrary components. Miller identified a mechanism comprising the type three secretory system (TTTS) used by parasitic bacteria for pumping toxic substances through cell walls. Since TTTS is tugging molecules through itself, it is a rudimentary version of the flagellar motor which tugs the molecules of the axle round and round. Thus, states Dawkins, evolution must have simply "commandeered" this component for the bacterial flagellum.

And so the game is revealed. Just what does "commandeer" mean? Perhaps evolution's "blind watchmaker," whom Dawkins sees working by "trial and error," is peeking under his blindfold. Did evolution devise the *programs* for the selection of the components, the fittings, and the modifications necessary? Then, as we have just seen, evolution must be employing a far more powerful "device" than a Turing class computer. Michael Shermer (2006) quotes Darwin's concept of "exaptation":

> On the same principle, if a man were to make a machine for some special purpose, but were to use old wheels, springs, and pulleys, only slightly altered, the whole machine, with all its parts, might be said to be specially contrived for that purpose. Thus throughout nature almost every part of each living being has probably served, in a slightly modified condition, for diverse purposes, and has acted in the living machinery of many ancient and distinct specific forms. (Darwin, quoted by Shermer, p. 68)

Though Darwin is clearly going to be no better off than Miller in coaching AI on the design of mousetraps, in lieu of "commandeer," Shermer confidently employs the term "co-opt," as in evolution "co-opts" features to use for another purpose. For "commandeer," Scott (2004) uses "borrowing and swapping." For "commandeer," Dennett (1996) substitutes the term "generate and test," holding, with no explication, that evolution simply "generates" new devices such as flagellar motors (or mousetrap #5) to test them out. Finally, Kevin Miller himself simply uses "mix and matching" saying, ". . . it's to be expected that the opportunism of evolutionary processes would mix and match proteins to produce new and novel functions" (2004, p. 88). If Dennett, Shermer, or the evolutionary biologists know secretly how to program these things, if they have solved the problem of commonsense knowledge, they should be teaching the folks in AI.

Programming in Evo Devo

Perhaps it may be felt that the recent discoveries of "Evo Devo" (Carroll, 2005) obviate these arguments. It is now understood that all complex animals — people, flies, trilobites, dinosaurs, and butterflies — share a common "tool kit" of master genes that govern the formation and patterning of their bodies and body parts. With this tool kit, fish fins can be modified into the legs of terrestrial vertebrates, or a simple tube-like leg can be modified into a wing. The development of these forms depends upon the turning on and off of genes at different times and places in the course of development, especially those genes that affect the number, shape, and size of a structure. Further, about 3% of our DNA or roughly 100 million bits is regulatory in nature. This DNA is organized into "switches" that integrate information about position in the embryo and the time of development.

In some essential respects, then, we have discovered a programming language. It is a language that interfaces with the concrete, biological world, and programmed correctly, can produce complex, concrete, functioning forms. But Freeman and Newell also, in their manipulation and matching of functional provisions of objects to functional requirements, fully intended this to be done in a programming language. As in any complex language, its effect (its semantics) depends entirely on the correct sequencing of its instructions. It must form a proper program — or it either "blows up" with logic errors or produces gibberish. Unless you wish to be ridiculed by the programming profession, the complex, programmed sequence does not happen by chance, no more than the instructions of a JAVA program to display a web screen occur by luck. Some one, some thing, some force guides the sequencing derived from the complex and rich instruction set and syntax available. A flick of a "switch" to the wrong value and a leg grows on top of a fly's head — or a useless spring is placed at the wrong position on the mousetrap.

The problem posed by Behe's humble mousetrap remains in full force. Nothing has changed. The use of a language *still implies knowledge of its semantics*, and in the mousetrap context, this still involves the transformations, positioning, fabrications, fittings, and fastenings of parts that all work toward a concrete function and which must enfold invariance laws. The smug rejection of mousetraps should cease, and the deep problem they represent be addressed. Until then, I expect that we still will see liberal use of the equivalents of "co-opting" and "commandeering," now appearing in statements such as "evolution created this new instruction set," or it "modified this instruction set."

This is not to mention one other obvious fact: there are many languages — JAVA, COBOL, FORTRAN, C++, Assembler, BASIC. I have yet to hear of one that was discovered just laying around, or that defined itself and published a user manual. Some one dreamt it up. If the powerful gene/switch language is an exception, how did this occur?

Conclusion

This discussion should not be construed as an argument for Intelligent Design in evolution. In *Creative Evolution*, with detailed argument, Bergson (1907/1911) rejected both radical mechanism and finalism. In radical mechanism we see the vision, accepted by Dennett and inherent in Darwin, of the great universal machine, unrolling or unfolding its forms and creatures, with deterministic precision. The word "time" means nothing to this conception. It has never taken to heart the implications of the simple fact that where time is melodic, where each "instant" is the reflection of the whole history of change — nothing can truly repeat. This undermines the very notion of deterministic causality.

Finalism is Bergson's term for the conception that the universe is the result of a vast plan, an enormous idea or conception. It is simply the inverse compliment of radical mechanism. Where radical mechanism drives towards the end result via its laws and initial conditions, finalism, from the other direction, draws the results irresistibly to the fulfillment of the great idea. The unforeseen creativity of real, concrete time is eclipsed. Finalism, too, cannot spell t-i-m-e, and Intelligent Design, particularly when taken "from the beginning of things," is in the end — finalism.

It was with deep thought that Bergson himself directed his own ship, steering a direction between finalism and mechanism. He held to a vision of evolution which respects the nature of time. His vision has been rejected in knee-jerk fashion as "vitalism," though in fact he critiqued the vitalist position. But perhaps we are nearing the point when a more profound direction of thought on evolution and time, and on mind and mousetraps, can be considered.

References

Behe, M. (1996). *Darwin's black box: The biochemical challenge to evolution*. New York: The Free Press.

Behe, M. (2000). A mousetrap defended: Response to critics. http://www.arn.org/docs/behe.mb_mousetrapdefended.htm

Behe, M. (2007). *The edge of evolution: The search for the limits of Darwinism*. New York: Free Press.

Bergson, H. (1911). *Creative evolution*. New York: Holt. (originally published 1907)

Bergson, H. (1912). *Matter and memory*. New York: Macmillan. (originally published 1896)

Carroll, S. (2005). *Endless forms most beautiful: The new science of Evo Devo and the making of the animal kingdom*. New York: Norton.

Chalmers, D.J., French, R.M., and Hofstadter, D. (1992). High level perception, representation and analogy: A critique of artificial intelligence methodology. *Journal of Experimental and Theoretical Artificial Intelligence, 4*(3), 185–211.

Copeland, B.J. (2000). Narrow versus wide mechanism: Including a re-examination of Turing's views on the mind–machine issue. *Journal of Philosophy, XCVI*, 1, 5–32.

Crowder, R.G. (1993). Systems and principles in memory theory: Another critique of pure memory. In A. Collins, S. Gathercole, M. Conway, and P. Morris (Eds.), *Theories of memory* (pp. 139–161). Hillsdale, New Jersey: Erlbaum.

Dawkins, R. (2006). *The God delusion*. New York: Houghton–Mifflin.

Dennett, D. (1996). *Darwin's dangerous idea: Evolution and the meanings of life*. New York: Simon and Schuster.

Dietrich, E., and Markman, A.B. (2000). Cognitive dynamics: Computation and representation regained. In E. Dietrich and A.B. Markman (Eds.), *Cognitive dynamics: Conceptual and representational change in humans and machines* (pp. 5–29). Hillsdale, New Jersey: Erlbaum.

Doumas, L., Hummel, J., and Sandhofer, C. (2008). A theory of the discovery and predication of relational concepts. *Psychological Review, 115*, 1–43.

Fodor, J., and Pylyshyn, Z. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition, 28,* 3–71.

Freeman, P., and Newell, A. (1971, September). A model for functional reasoning in design. *Second international conference on artificial intelligence*. London.

French, R.M. (1990). Sub-cognition and the limits of the Turing test. *Mind, 99*, 53–65.

French, R.M. (1999). When coffee cups are like old elephants, or why representation modules don't make sense. In A. Riegler, M. Peshl, and A. von Stein (Eds.), *Proceedings of the international conference New Trends in Cognitive Science* (pp. 158–173). Austrian Society for Cognitive Science.
Gelernter, D. (1994). *The muse in the machine: Computerizing the poetry of human thought*. New York: Free Press.
Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science, 7,* 155–170.
Gibson, J.J. (1966). *The senses considered as perceptual systems*. Boston: Houghton–Mifflin.
Goldinger, S. (1998). Echoes of echoes? An episodic theory of lexical access. *Psychological Review, 105*, 251–279.
Hanson, N.R. (1958). *Patterns of discovery*. Cambridge: Cambridge University Press.
Haselager, W. (2005). Robotics, philosophy and the problem of autonomy. *Pragmatics and Cognition, 13,* 515–532.
Holyoak, K.J., and Thagard, P. (1997). The analogical mind. *American Psychologist, 52*, 35–44.
Hummel, J.E., and Holyhoak, K.J. (2005). Relational reasoning in a neurally plausible architecture. *Current Directions in Psychological Science, 14,* 153–157.
Ingerman, P. (1966). *A syntax-oriented translator*. New York: Academic Press.
Isaak, M. (1997). Bombardier beetles and the argument of design. http://www.talkorigins.org/faqs/bombardier.html
Katz, J.J., and Fodor, J.A. (1963). The structure of a semantic theory. *Language, 39*, 170–210.
Kim, N., Turvey, M., and Carello, C. (1993). Optimal information about the severity of upcoming contacts. *Journal of Experimental Psychology: Human Perception and Performance, 19,* 179–193.
Klein, D.B. (1970). *A history of scientific psychology*. New York: Basic Books.
Kofahl, R.E. (1981). The bombardier beetle shoots back. *Creation/Evolution, 2*, 12–14.
Kugler, P., and Turvey, M. (1987). *Information, natural law, and the self-assembly of rhythmic movement*. Hillsdale, New Jersey: Erlbaum.
Lieber, L., and Lieber, H. (1945). *Einstein's theory of relativity*. New York: Holt, Rinehart, and Winston.
Lloyd, S. (2006). *Programming the universe*. New York: Alfred A. Knopf.
Lynds, P. (2003). Time and classical and quantum mechanics: Indeterminacy versus discontinuity. *Foundations of Physics Letters, 16,* 343–355.
McCarthy, J., and Hayes, P. (1969). Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer and D. Michie (Eds.), *Machine intelligence* (Vol. 4, pp. 463–502). Edinburgh: Edinburgh University Press.
McDonald, J. (2000). A reducibly complex mousetrap. //http://udel.edu/~mcdonald.oldmousetrap.html
McDonald, J. (2002). A reducibly complex mousetrap. //http://udel.edu/~mcdonald.mousetrap.html
Melendez–Hevia, Waddell, T., and Cascante, M. (1996). The puzzle of the Krebs citric acid cycle: Assembling the pieces of chemically feasible reactions, and opportunism in the design of metabolic pathways during evolution. *Journal of Molecular Evolution, 46,* 508–520.
Miller, K. (2003). Answering the biochemical argument from design. In N. Manson (Ed.), *God and design: The teleological argument and modern science* (pp. 292–307). London: Routledge.
Miller, K. (2004). The flagellum unspun: The collapse of irreducible complexity. In M. Ruse and W. Dembski (Eds.), *Debating design: From Darwin to DNA* (pp. 81–97). Cambridge: Cambridge University Press.
Nottale, L. (1996). Scale relativity and fractal space-time: Applications to quantum physics, cosmology and chaotic systems. *Chaos, Solitons and Fractals, 7,* 877–938.
Paivio, A. (1971). *Imagery and verbal processes*. New York: Holt, Rinehart, and Winston.
Penrose, R. (1994). *Shadows of the mind*. Oxford: Oxford University Press.
Prinz, J., and Barsalou, L. (2000). Steering a course for embodied representation. In E. Dietrich and A.B. Markman (Eds.), *Cognitive dynamics: Conceptual and representational change in humans and machines* (pp. 51–77). Hillsdale, New Jersey: Erlbaum.
Pylyshyn, Z. (1973). What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin, 80,* 1–22.
Pylyshyn, Z. (2002). Mental imagery: In search of a theory. *Behavioral and Brain Sciences, 25,* 157–237.
Robbins, S.E. (1976). Time and memory: The basis for a semantic-directed processor and its meaning for education. Doctoral dissertation, University of Minnesota, Department of the Psychological Foundations of Education.

Robbins, S.E. (2000). Bergson, perception and Gibson. *Journal of Consciousness Studies, 7,* 23–45.

Robbins, S.E. (2001). Bergson's virtual action. In A. Riegler, M. Peschl, K. Edlinger, and G. Fleck (Eds.), *Virtual reality: Philosophical issues, cognitive foundations, technological implications* (pp. 189–202). Frankfurt: Peter Lang Verlag.

Robbins, S.E. (2002). Semantics, experience and time. *Cognitive Systems Research, 3,* 301–335.

Robbins, S.E. (2004a). On time, memory and dynamic form. *Consciousness and Cognition, 13,* 762–788.

Robbins, S.E. (2004b). Virtual action: O'Regan and Noë meet Bergson. *Behavioral and Brain Sciences, 27*, 906–907.

Robbins, S.E. (2006a). On the possibility of direct memory. In V.W. Fallio (Ed.), *New developments in consciousness research* (pp. 1–64). New York: Nova Publishers.

Robbins, S.E. (2006b). Bergson and the holographic theory of mind. *Phenomenology and the Cognitive Sciences, 5,* 365–394.

Robbins, S.E. (2007). Time, form and the limits of qualia. *Journal of Mind and Behavior, 28*, 19–43.

Robbins, S.E. (2008). Semantic redintegration: Ecological invariance. Commentary on T. Rogers and J. McClelland (2008). Précis on *Semantic Cognition: A Parallel Distributed Processing Approach. Behavioral and Brain Sciences, 31,* 726–727.

Robbins, S.E. (2009). The COST of explicit memory. *Phenomenology and the Cognitive Sciences, 8,* 33–66.

Robbins, S.E. (2010a). [Book review] *The case for qualia. Journal of Mind and Behavior, 31,* 141–156.

Robbins, S.E. (2010b). Special relativity and perception: The singular time of psychology and physics. *Journal of Consciousness Exploration and Research, 1,* 500–531.

Robbins, S.E. (2012). *Time and memory: A primer on the scientific mysticism of consciousness*. Charleston, South Carolina: CreateSpace.

Rogers, T., and McClelland, J. (2004). *Semantic cognition: A parallel distributed processing approach*. Cambridge, Massachusetts: MIT Press.

Rogers, T., and McClelland, J. (2008). [Precis of] *Semantic cognition: A parallel distributed processing approach. Behavioral and Brain Sciences, 31*, 689–749.

Savelsbergh, G.J.P., Whiting, H.T., and Bootsma, R.J. (1991). Grasping tau. *Journal of Experimental Psychology: Human Perception and Performance, 17,* 315–322.

Scott, E. (2004). *Evolution vs. creationism*. Berkeley: University of California Press.

Shapiro, J. (1996, September). In the details . . . what? *National Review,* 62–65.

Shepard, R.N., and Feng, C. (1972). A chronometric study of mental paper folding. *Cognitive Psychology, 3,* 228–243.

Shermer, M. (2006). *Why Darwin matters: The case against intelligent design*. New York: Times Books.

Sloman, A. (1971, September). Interactions between philosophy and artificial intelligence: The role of intuition and non-logical reasoning in intelligence. *Second international conference on artificial intelligence.* London.

Turvey, M., and Carello, C. (1995). Dynamic touch. In W. Epstein and S. Rogers (Eds.), *Perception of space and motion* (pp. 27–66). San Diego: Academic Press.

Wertheimer, M. (1945). *Productive thinking*. New York: Harper and Row.

Wigner, E.P. (1970). *Symmetries and reflections*. Cambridge, Massachusetts: MIT Press.

Woit, P. (2006). *Not even wrong: The failure of string theory and the search for unity in physical law.* New York: Basic Books.

Wolff, C. (2010). *Psychologia empirica methodo scientifica pertractata* . . . . Whitefish, Montana: Kessinger Publishing. (originally published 1732)

Woodward, J. (2000). Explanation and invariance in the special sciences. *British Journal for the Philosophy of Science, 51,* 197–214.

Woodward, J. (2001). Law and explanation in biology: Invariance is the kind of stability that matters. *Philosophy of Science, 68,* 1–20.

Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford: Oxford University Press.

Young, M., and Edis, T. (2004). *Why intelligent design fails*. New Brunswick, New Jersey: Rutgers University Press.