ARTICLE IN PRESS

ELSEVIER
Cognitive Systems Research 1 (2002) 000–000
Cognitive Systems RESEARCH
www.elsevier.com/locate/cogsys

Semantics, experience and time

Action editor: Ron Sun

Stephen E. Robbins
Center for Advanced Product Engineering, Metavante Corporation, 10850 West Park Place, Milwaukee, WI 53224, USA

Received 27 April 2001; accepted 5 November 2001

Abstract

The computational hypothesis, with its inherent representationalism, and the dynamical hypothesis, with its apparent absence of representations and its commitment to continuous time, stand at an impasse. It is unclear how the dynamical stance can handle representational thought, or how computationalism can deal effectively with a tightly coupled, reciprocally causative agent–environment system in continuous transformation. Underlying this dilemma is the complex relation of mind to time, a relation encoded in the word *experience*. We must ask if any hypothesis describes a 'device' capable of experience? Yet what is an intelligence and its thought without experience? Is a computational device, whether supporting a symbolic processor or connectionist net, intrinsically condemned to a zero degree of experience? What is required of a dynamical device? It is argued here that 'semantic' intelligence and thought rests upon experience, fundamentally upon the invariance laws defined over time within conscious perception. The structure of experience is intrinsically unavailable to the computational device, limiting it to a 'syntactic' intelligence. An alternative conception of a device is offered, based on Bergson conjoined with Gibson, which supports the qualitative and structural aspects of experience and the semantic. It frames a dynamical model of perception and memory in which invariance laws are intrinsic, creates a deeper notion of situatedness, and supports a concept of semantically based, representative thought founded upon perception. © 2002 Published by Elsevier Science B.V.

Keywords: Semantics; Experience; Time

1. Introduction

Let us imagine a dynamic system, and with it, the rich representational world it must express. Consider 'paddling,' as in paddling a standard canoe with a single-blade paddle. We shall do this from the first person perspective, as though we actually were at the stern. There is first of all the visual experience. Always present are the side and form of the canoe, directed at some point far down the lake, and up front, the bow paddler (hopefully) also in motion. There is the expanse of water, usually flecked and rippled by the wind—a texture field and gradient after Gibson's heart. This whole field is flowing, moving by the canoe—a classic optical flow field centered largely on the distant portage point—the focus of optical expansion. Enriching the texture may be light-sparks, dancing, from the sunlight. Into this moving fluid, a wooden blade is plunging, pulling straight back, leaving a boiling gap in the

---

E-mail address: stephen.robbins@metavante.com (S.E. Robbins).

1389-0417/02/$ – see front matter © 2002 Published by Elsevier Science B.V.
PII: S1389-0417(02)00045-1

surface. The blade lifts from the water, feathers, and returns in an arc, trailing a precise pattern of droplets splashing in the water as it returns to its starting point. There are multiple acoustical invariants in this, but we isolate one—the 'thwunking' sound of the paddle plunging and pulling through the fluid. The haptic experience of this 'wielding' of the paddle is partially described by the inertial tensor or invariant, I_{ij}, described so exhaustively by Turvey (cf. Turvey & Carello, 1995), though we may augment this description with the strain on the left shoulder as it drives the paddle, the strain and tension on the right arm as it pulls and swings, and, since we have a larger system than the wielding of the paddle, there is the response and thrust of the canoe and its simultaneous inertial resistance and felt mass, the felt pressure of the wind, the force of the water. Paddler–agent and environment are indeed tightly coupled in this system, locked in a reciprocally causal relationship. The fluid medium provides instant, continuous feedback to the paddle relative to its force and direction. The mass, inertia and form of the canoe relative to the surface and to the point of optical expansion, even the sensed change of wind on the face, provide continuous and instant feedback on the effect of the stroke, causing further adjustments—the slight flick of the blade via a j-stroke or a sweep-stroke to correct the line.

For this multi-modal system, dynamically flowing in time, with its reciprocally causal feedback loops, with its invariance defined across modalities and over time, there are two primary contenders for explanation today. These are the computational hypothesis (CH), with its commitment to mental representations, and the dynamical hypothesis (DH) (Port & van Gelder, 1995; van Gelder, 1998), admittedly in its infancy, yet posed as a contender precisely because it is more amenable to a natural dynamics. But these two are locked at an impasse. There is a fundamental intuition behind representations, namely that we as cognitive agents are able to represent (or imagine) an environment and plan action in the absence of its immediate presence. Our paddler may have planned his route and goal before he got into the canoe or even looked at the lake. Yet this capacity does not seem naturally addressable by the DH. But on the other side, the continuous, in time, system-as-a-whole nature of perception/action appears easily addressable only by the DH.

1.1. Representation vs. dynamics

The reciprocal causation and tight coupling of the DH (Clark, 1997) generate a difficulty for computationalism, already glimpsed above, namely a form of boundary problem (cf. Grush, 1997b). If agent–environment truly form a single, dynamic system, can we truly draw a boundary between some inner, cognitive space, and a non-cognitive, passive environment? This is both a spatial and a temporal question. The feedback loops in the system are continuously and temporally extended. The system resists a simple decomposition to temporal components on the agent side and represented, environmental, temporal elements. The evolutionary equations that describe such a coupled system incorporate terms that factor in the current states of both components, and the system is treated as an evolving whole.

This stance, which at face value appears to exclude representations in its explanatory framework, has been challenged for precisely this reason as being incomplete. The DH is perhaps adequate for a merely *adaptive* system. A cognitive agent, it is argued, requires representations precisely because it can plan or reason about environments with which it is not immediately in contact. While we would hesitate to call our canoe paddler a non-cognitive agent simply because he is mindlessly, adaptively paddling down the lake (cf. Chopra, 1999), he does have a goal, and we can take the point that the paddler indeed engages at times in less purely adaptive, more representative or planning modes. Clark (1997) and Grush (1997a,b), among others, have attempted to augment the DH with representational apparatus. The (dynamical) agent (A) is thus fitted out with, (a) a capacity to 'emulate' the environment (E) for use as an internal representation, and (b) a 'controller' that decouples A from E, coupling instead to the environmental representation E′. The E↔A relation as a presentation (or perception) is considered tightly coupled, reciprocally causative, and 'weakly represented.' The A↔E′ relation is uncoupled, strongly representational.

This 'emulation' augmentation is a good index of the impasse. The emulative notion of E′ is a search for a small home for representationalism within the dynamical brain. In what form an actual emulation of E (lake, flow field, sun sparks, flashing paddles, etc.) could exist, or how the DH could actually grant this haven, given we see the brain as *globally* dynamical, is unclear. Meanwhile, the augmentation position has virtually surrendered *perception* to the DH. Within the reciprocally causal, coupled relation of E↔A, it reserves some 'weakly representational' feature, but why the DH should (or could) even grant this is also unclear. E′ attempts to preserve the role of imagination and/or cognition for the representational approach. But what if cognition is built upon perception, i.e., what if it essentially involves more complex forms of the same dynamics?

The augmentation solution is driven by multiple frustrations. There is first of all the difficult status of the descriptive entities of the DH—attractors, bifurcations, trajectories—as 'representations' of the world within the brain. Even granting perception to the DH, there is no consensus, even by DH theorists, as to how these constructs are ever translated into the world as phenomenally experienced. 'Standing-in for' (Bechtel, 1998) buys us little here. The DH can insist on its tight coupling of agent and environment, but how then is this fact effectively used in explaining perception vis a vis these fundamental constructs of the DH? The DH in fact maximizes the explanatory gap—the description of the happenings in the brain versus our perceptual experience of the world-out-there while paddling down the lake is now openly maximal. Only the deceptively clear refuge of representations (whose true definition has grown extremely vague) appears to give relief. But in perception, representations never escape the homunculus. They simply become a code that some homuncular eye must unfold as the perceived world. In their retreat to the world of cognition and imagination, one must ask if they do any better, now being 'unfolded' as a mental image.

Somehow the constructs of the DH must be shown to have true utility in the explanation of perception, i.e., in truth, the constructs of the DH must be used to leap the infamous explanatory gap. Simultaneously, the framework employed must embrace the fundamental intuition behind representationalism, i.e. it must support the universal experience of *representative* thought.

1.2. Semantics, experience and time

In the paddle up the lake, we were describing an *experience*. Indeed, as French (1990) argued in the context of the Turing Test, it is our natural experience of the every day world that is so critical and difficult to duplicate in any computational representation, and it is questionable, he thought, whether any computer in the foreseeable future, lacking the ability to experience, can ever pass the Test. Here, in the Test, representation and semantics inevitably meet.

The Turing Test is fundamentally a problem of semantics. The very form of the Test reflects our faith that language is a reflection of mind. But the problem of semantics, with respect to language, and precisely the linguistically embodied questions asked in the Test, becomes the question of how the discrete symbols of a language evoke the *experience* of a mind. The question becomes whether a computing machine, be it even a sensory-equipped mobile robot harboring a connectionist net, or harboring programs in the symbolic paradigm, or both, is in principle capable of experience, i.e., the experience referenced by a language? Could it ever have the knowledge we have after one paddle up the lake? What is it that might in truth forever condemn the computing machine (as currently defined) to experience degree zero? And what is required of even a non-computational (if so construed) dynamical system such as Brooks' (1991) robot?

I do not speak here of 'experience' as it is defined in machine learning, i.e., the exposure of a program to samples of information over time, feedback processes and/or evolutionary self-modification of programs or weights. 'Experience,' as French used the term, is fundamentally *perception*, and perception is intrinsically over time. It is the intrinsic semantics of perception that is primary. The flow field of the lake as the paddler strokes is intrinsically meaningful. The dynamic E↔A relationship is itself profoundly semantic. A device that cannot support and retain in memory (as A↔E′) this meaningful perceptual flow we call experience, it can be argued, will never pass the Test.

It is the relation of mind to time that makes this experience so problematic for our theories. To solve it, the DH must embrace the tightly coupled pair of agent/environment in a far more profound way than is currently understood. To paraphrase Port and van Gelder (1995), it must get far more serious about this coupling. Though they noted in their introductory discussion of the DH that we must get serious about *time*, the DH itself must get more serious about time. In truth, it must incorporate this statement of Bergson, a statement we will revisit: ‘‘Questions relating to subject and object, to their distinction and their union, must be put in terms of time rather than of space’’ (1896/1912, p. 77).

We will explore then a route out of the impasse. The dynamical hypothesis can be placed within the framework created when the theories of Henri Bergson (1896/1912), with his vision of time, and J.J. Gibson (1950, 1966, 1979), with his notion of the role of invariance, are joined. From this emerges a basis for supporting experience, semantics and the fundamental intuition behind representations—representative thought.

2. The dynamic structure of time-extended experience

It is my purpose here initially to paddle us through the problem of experience and its time-extension. We must look more closely at what we are trying to explain. It is common in the literature on the CH side (e.g., Prinz & Barsalou, 2000; Dietrich & Markman, 2000) to hear that discrete symbols, riding on top of a continuous dynamics, can support experience and semantics. The problem has not been fully understood. We must truly understand what we are trying to ‘represent.’

Firstly, we must focus on the invariance laws which define our experience of time-extended events. These are essential to our understanding of the dynamic structure of experience. As an initial context, let us consider a Turing Test question French (1990) proposed to apply to some hopeful human-imposter of a machine. This ‘subcognitive’ question, as well as others he applied, fundamentally rests on the perception of invariance. Thus one of the main phenomena that French relied upon is associative priming. Presented the word ‘bread,’ it takes a human significantly less time to recognize ‘butter’ if this is the word following as compared to a word less closely ‘associated’ such as ‘vase.’ French’s hypothetical interrogator, armed with human response time norms for target words and their various associates, was to require the machine to produce the same response time pattern. But these norms, French pointed out, come from concrete human experience. How could the mass of associative strengths for all possible pairs of words be pre-programmed?

French noted that the associative strengths of these concepts often come from their close association in a sequence. The steps in baking a cake—finding the box in the cupboard, getting the bowl, stirring the batter, opening the oven—are all associated by contiguity of experience. The are innumerable such groups—swimming, going to the store, canoeing, eating breakfast, etc.—each with a temporal order which the associative strengths reflect. He noted that though it might be possible theoretically to program these a priori, the possibility is dim. To establish the weights, the machine would in reality require the experiences.

Initially, this might sound vaguely like something for which a connectionist device would be more amenable. But let us look more closely.

2.1. Multi-modal invariance

We will consider closely an event such as ‘stirring,’ as in stirring a cup of coffee. There are multiple Gibsonian (1966, 1979) invariants defined over time and over the various modalities of this event. I note initially that we can speak of two forms of invariants—structural and transformational (Shaw & Wilson, 1974). By ‘transformation’ is meant that information specific to the ‘style’ of a change, e.g. the information defining bouncing, rolling, rotating, expansion, contraction, opening, swirling, etc. By ‘structural invariant’ is meant that information specific to the thing or object undergoing the change, e.g. the ball, the balloon, the dough, the cube, the coffee. For the (slow) ‘event’ of the aging of the facial profile (Pittenger & Shaw, 1975), the structural invariant is abstractly defined mathematically as a

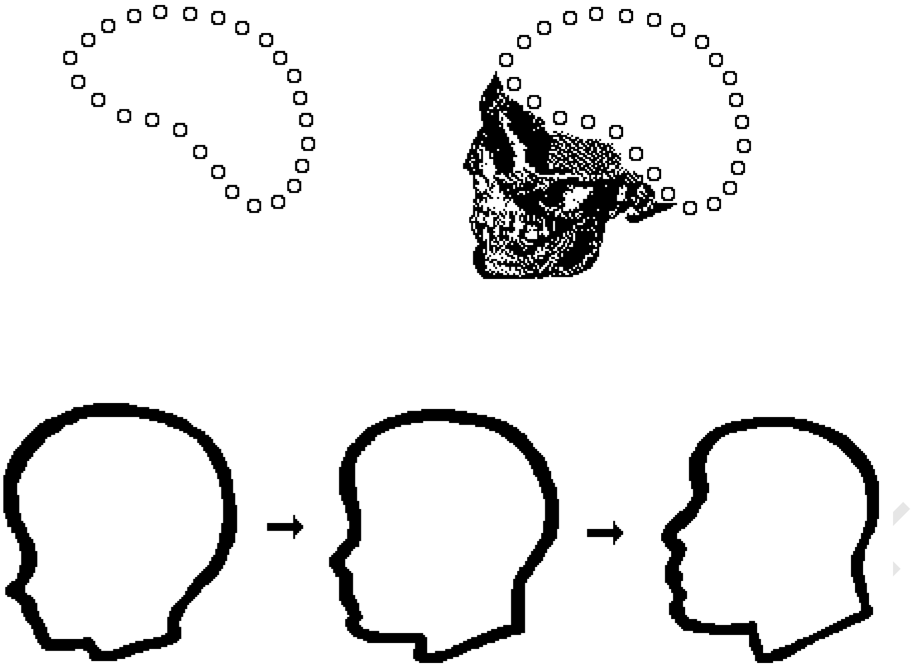

Fig. 1. Aging of the facial profile. A cardioid is fitted to the skull and a strain transformation is applied. (Strain is equivalent to the stretching of the meshes of a coordinate system in all directions.) Shown are a few of the possible profiles generated. (Adapted from Pittenger & Shaw, 1975.)

cardioid (Fig. 1). The transformation which grows and ages the profile is defined mathematically as strain (equivalent to stretching the coordinate system on which the cardioid is placed). In the case of an optical flow field, as created for example when we drive down a road, we have the lawful expansion of the field according to the relation $V \propto 1/D^2$, where this relation defines the inverse relation of the values the velocity vectors to the distance from the observer (Fig. 2). This is the transformational information, while there is information specific to the thing undergoing this transformation, e.g. a texture gradient specific to a field of grass (Fig. 3).

For our coffee stirring event, there is the visual invariance—a radial flow field centered on and radiating from the stirring object. There is the auditory invariance—the steady clinking sound of metal spoon against cup. There is the olfactory invariant—a certain coffee aroma. In the dynamics of the haptic component of the event, we can define the motion by the wielding of a 'tensor object' that captures the inertia tensor (invariant), I_{ij}, specific to spoon-stirring (cf. Turvey & Carello, 1995). Gelernter (1994) envisioned an operation of taking a 'stack' of these events/experiences such that invariants are defined across the stack while variants 'wash out.' In the 'stirring' case, for example, there are invariants

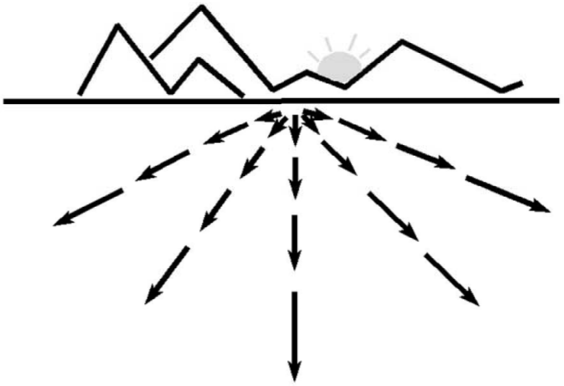

Fig. 2. Optical flow field with gradient of velocity vectors.

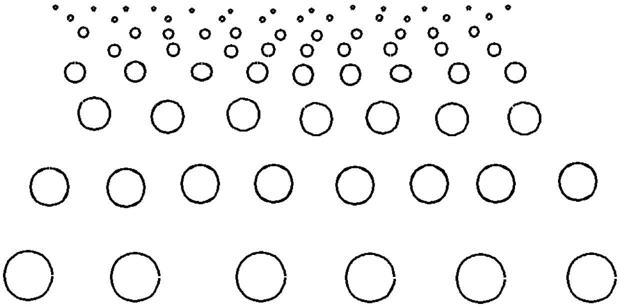

Fig. 3. Texture density gradient (Gibson, 1950). The horizontal separation, *S*, is proportional to the distance, or $S \propto 1/D$, the vertical separation as $S \propto 1/D^2$.

such as a liquid medium being and capable of being stirred, an instrument with sufficient structural rigidity and width to move the liquid, a source of force to move the instrument, a container spatially constraining the liquid. These are equally invariance laws defining this event. This is a second sense of the term 'invariant,' used here in the context of, or with respect to, this operation of 'stacking' of events. Later (Section 3.5), a more concrete mechanism for this operation will be discussed.

All the elements of this invariance structure are naturally, intrinsically 'associated,' but are so precisely due to their structural roles in the event. The spoon, in this coffee stirring event, is an invariant of 'normal context,' and it fulfills the required instrument with structural properties needed for moving the medium. 'Cup' is also an invariant, fulfilling the requirement for spatially constraining the liquid. 'Spoon' is an 'associate' of 'coffee,' as is 'cup,' as is 'stirring,' etc., as is the haptic form, as is the visual form, as is the sound for which we have no particular name but which must equally be an 'associate.' How else to explain the anomalous feeling presented by a sentence such as, ''The coffee crackled and popped as he stirred?'' This is a violation of the auditory invariance of the event. Or, for other violations of nameless 'associates,' i.e., invariants:

- ''As he stirred, the coffee gushed small geysers of liquid.''
- ''As he stirred, the spoon slowly bent and collapsed.''
- ''As he stirred the coffee, the formaldehyde aroma wafted from the cup across the room.''
- ''Drawing a string of spaghetti from his plate, he stirred the coffee.''
- ''As he stirred the coffee, his thigh muscles quickly fatigued.''

For all of these violations, we can create a context in which they make sense. An evil lab-assistant has poured formaldehyde into a fellow technician's coffee. The collapsing spoon has been weakened by previous super-heating, or the coffee has been pressure-heated to 400 degrees, or the spoon is cheap plastic. The spaghetti is very thick and the 'stirring' is extremely half-hearted, barely qualifying as such. Well, maybe not for the 'geysers,' but this kind of violation of invariance is precisely what makes cartoons funny. The point is that understanding even these 'anomalous' sentences (events) yet rests on the invariance structure. The 'context' adjustments are changes in global parameters that allow modifications to the local invariance. The spoon collapse invokes an invariance law of heat relative to plastic under which the stirring event is now perceived. Suppose the event:

- Rolling down the hill, the styrofoam rock crushed the brick house.

As context, I say ''Japanese monster movies.'' It is the invariance structure of 'crushing' that is key. Houses in normal context have rigidity properties that require a certain force to preserve the invariance defining a 'crushing.' The context specification changes these globally in an instant; the invariance structure holds.

These structures are *events*—inherently multi-modal structures of dynamic transformations and invariance over time. If we consider representing these as weights in a connectionist net, we see that we now have a homogeneous medium and a static representation. The multi-modal aspect is gone, and the continuous transformations preserving invariance are gone as well. Though it is not uncommon (as did French, 1990) to reference the standard, computer friendly 'spreading activation' model to explain the priming effect of, for example, 'spoon,' this model likely has an illusory validity. The 'associates' are unlikely to be 'nodes' in a network connected by 'strengths.' The multi-modal, dynamical event with its many nameless invariants would be extremely difficult to so statically and homogeneously represent. A dynamic system offers at least a partial solution, though the multi-modality is yet problematic. The dynamic transformations of neural patterns defined globally over the brain and supporting the perception of this event across modalities would be described as a form of attractor. We could then think of presenting the word 'spoon' as re-invoking the complex dynamical pattern defined globally over the brain and supporting stirring's invariance structure. This difference would indicate why the spreading activation model is silent in the face of context sensitive effects. Relying on normal context, we present 'spoon'—'coffee' is primed. We preface the presentation with a sentence, ''Halloween is a fun night.'' Now presentation of 'spoon' primes 'pumpkin,' as in spooning out a pumpkin's insides. Did all the stored pair-wise 'strengths' suddenly change? Or did the new context bias a dynamic system towards the evocation of an entirely different dynamic pattern (cf. Klinger & Burton, 2000)?

But to the primary point, we have had an initial view of the invariance structure of events, i.e., of experience, and glimpsed how the basis for *semantic* understanding, even of a simple sentence like, ''The man stirred the coffee,'' must be founded on an invariance structure. As we shall see, the most complex syntactic relations of abstract symbols cannot do justice to these multi-modal and dynamical transformations.

2.2. Event/object features as invariants

French also proposed a 'rating test' for the computer, again comparing the responses to human norms. For example:

- ''Rate *banana splits* as *medicine*''
- ''Rate *pens* as *weapons*''
- ''Rate *jackets* as *blankets*''
- ''Rate *purses* as *weapons*''

These, he argued, involve the overlap of two categories, and again, he argued, it is ''virtually impossible'' to explicitly program all the degrees and types of ''associations'' required to answer the questions as would a human.

But this is not the overlap of two categories, particularly, as we shall see, if we take this 'overlap' to be the intersection of two sets or vectors of features, for example features of purses and of weapons respectively. Nor is it static 'associations.' It is *the projection of the transformational dynamics of an invariance structure upon a possible component*. Consider this rating task:

- ''Rate *knives* as *spoons*''

If we invoke, as context, the invariance structure of 'stirring' coffee, then under this transformation, the knife displays the requisite structural invariance to move the liquid—and thus rate quite well. If we invoke, as context, 'eating soup,' then under this transformation, the knife rates poorly. Suppose then the rating task:

- ''Rate *ducks* as *spoons*''

Under the coffee stirring transformation, the duck's bill proves to have the needed structural invariance. Do we hold then that ducks and spoons turn out to have had a nice 'association' all along? It should be clear from these few examples alone that not only is it a virtual impossibility to pre-program (or use predefined vectors for) these 'associations,' it is a *total* impossibility. It is a total impossibility because there is no static, finite, pre-defined set of object 'features' that can be compared and thus 'strength' related, even assuming the other impossibility—that one could store all possible comparisons. A 'feature' is simply an invariant under some transformation. Thus:

- "Rate *socks* as *fly-swatters*"

Under the transformation contemplated, the 'rigidity' feature of socks required for swatting suddenly emerges. It did not pre-exist. It would not have been pre-stored in a data structure. But this is reciprocal. There is equally no finite, pre-existing set of transformations.

There can be no such things then as categories "composed of many tiny (subcognitive) parts that can overlap . . . " (p. 64) as French hoped could (ultimately) account for these associations, given this is another name for the 'feature' approach. Nor can Lakoff (1987) and Johnson (1987) be invoked, as he hoped, as supporting this approach. It is not a difficult exercise to show that their image schemas fundamentally are invariance structures as well, and therefore more dynamic than French appeared to be considering (at least in 1990). To take one example, the fundamental schema of 'containership,' they argued, is experienced as things going into and out of the body, things going into and out of the visual field, and things going into and out of things in the visual world. In other words, it is a higher order of dynamic invariance existing across multiple invariance structures. We can then:

- "Rate *boots* as *vases*"

The transformation of inserting a bunch of flowers reveals the structural invariance of containership for a passable vase.

But if the *experience* that French found so problematic rests upon invariance, then we shall see that the relationship of invariance to time is even more problematic.

2.3. Invariance and time

Consider a wire cube in a darkened room, rotating slowly around a rod placed through the center of two opposing faces, and strobed periodically in phase with or at an integral multiple of its symmetry period. The information specifying the shape of the cube is carried, over time, by this symmetry period. In this case, since it maps onto itself every 90°, a period of four. If this information is destroyed, e.g. if the cube is strobed arhythmically, it becomes a distorted, wobbly figure (Turvey, 1977b). This is clearly an invariance (symmetry) specified over time. The resultant of the arhythmic strobing implies that an arbitrary sampling rate (or set of discrete samples) fails to preserve this transformation and the structural invariant defined over time. The sampler, at the least, would have to be pre-adjusted to the rate, but what if, as Turvey noted, there were two cubes rotating at different rates, etc? But now let us consider a normal cube as it rotates, and gradually increase the velocity of rotation. We see that the cube transitions through a series of figures with increasing numbers of serrated edges—8, 12, 16 . . . , each an integral multiple ($4n$) of its symmetry period. Finally, at a high enough rate, it becomes a cylinder surrounded by a fuzzy haze, i.e., a figure of infinite symmetry.

Supporting these perceptual transitions, we can again posit an attractor supported over the transforming neural patterns of the brain. The attractor must be 'specific,' to use Gibson's term, to the form of the cube as it transforms. There is not an *instantaneous* cross-section of time (or point in the phase space) that captures the invariance specific to the cube. The invariant is not a 'bit' of information that can be transmitted along the nerves. The invariance exists only over time. The information specified *over time* can be destroyed in the case of the arhythmically strobed cube. But there is a deeper point.

Dynamical systems are systems that naturally *integrate* scales. The combined action of a myriad of smaller scale elements forms a large scale pattern. As we apply heat to the bottom of our coffee cup, or more precisely, something like Libchaber's (Gleick, 1987) fluid container, the number of cylindrical rolls of fluid, as described (initially) by the Lorenz attractor, continuously increase. Thus actions of a myriad of coffee molecules are coordinated to form large scale 'rolls.' Similarly, in the body/brain, there is a nested hierarchy of scales (cf. Keijzer, 1998), each level being inclusive of the next. The actions of myriads of atomic elements form large scale molecular movements. The action of myriads of neurons form large scale neural patterns. It is this hierarchical dynamics, we must assume, that determines the *time-scale* of the perceived world.

We perceive at a certain scale of time. The cube, rotating at a certain velocity and perceived as a figure with 16 serrated edges, is a perception relative to a certain scale of time. The fly buzzing by, his wings a-blur, is an index of our scale of time. This must be determined by the hierarchical dynamics of the brain. If we consider the brain, considered for a moment as simply a piece of the universal field, we see at the depths of this hierarchy, as physics tells us, 'particles' with life spans on the order 10^{-9} s and even vastly less. This is an incredibly rapid scale. From this we build to the slightly less rapid scale of quarks, then to the electrons, then to the molecular, then the neural. The total dynamics defined over these scalar levels determines our normal perceived scale. But at least in principle, it has been argued (e.g., Hoaglund, 1966; Fischer, 1966), this dynamics can be changed. We can introduce a change at a given level, but the system is a whole, and will be affected as a whole—there are no preferred 'levels' in a coherent system (cf. Ho, 1998). If we introduce a catalyst at the chemical level that modulates the orienting of appropriate bonds such that the velocity of chemical processes is increased, there will be an effect on the global dynamics. And there must be a perceptual consequence. The time-scale of the perceived world must change. Given a certain strength of catalyst, the fly may now be moving slowly by, his wings flapping like a heron's. The 16-edged cylinder-cube is now perceived as a four-sided cubical figure slowly rotating.

To borrow from physics, we will have changed the 'space–time partition.' And as in the physical theory, it is only invariance laws (e.g., $d = vt$, or $d' = vt'$) that hold across these partitions. The cube remains a figure of 4*n*-fold symmetry across partitions. The fly is specified by the same laws whether barely moving, or buzzing by. The aging of the facial profile, defined by its strain transformation applied to a cardioid figure, is specified by the same law across partitions, whether it becomes a fast event or an even slower event.

'Experience' then is intrinsically related to a *scale* of time, i.e., a given space–time partition, and the invariance laws defined over this partition which in turn specify events. We can begin to see then why the computing machine—if it lives in a scale-less world—would be at degree zero of experience. But it is the definition of scale, as we shall see more fully, that fundamentally supports the *qualitative* aspect of the perceived world.

2.4. Quality and time

Scale implies quality. The buzzing fly perceived at our normal scale, his wing-beats a-blur, is a certain quality. At the heron-like scale, it is a different quality. The color red, a proportion over trillions of oscillations of a field for but a second, is a certain quality. At a higher degree of the velocity of processes, where perception is closer to each developing oscillation, we have another, perhaps more vibrant quality of red. But *scale implies extent*. The dynamical state of the brain is specific (or proportional) to a given 4-D extent of time, i.e., to a set of past states of the universal field in which it (the brain) is embedded. The buzzing fly, as opposed to the heron-fly, represents a far higher ratio of events at the highest scale of the brain or organism (O) to events in the environmental field (E)—a proportion relative to a far greater history of events in the environment. As we raise the velocity of processes, the ratio (E/O) of events in the external field relative to events at the highest scale or 'level' of the brain lowers. The extent of the past specified in the heron-like case is far less than in the buzzing fly.

We cannot treat this extent as a series of discrete 'instants,' i.e., as a series of discontinuous 'states.' In doing so, the qualitative aspect of the fly's flight is destroyed, but to see this, suppose we were to do so, as though the motion of the fly were treated as a series of instantaneous 'snapshots.' What scale are these snapshots? We have no natural or normal scale on the world to invoke as is provided by the brain—we cannot stop our choice until we have plunged to the depths of the micro-scale of the universal field. Suppose we (arbitrarily) stop at 10^{-9} ns for the duration of each snapshot. Treating the motion of time this way as sets of 'present' instants or snapshots, each of which becomes instantly 'past', we now force the brain to 'store' each snapshot. From this enormous set, it must, in some totally non-understood manner, reconstruct the motion of the fly as a ‘composite’ and *specific to a given scale*, e.g., as the heron-fly. This treatment is invalid for at least two reasons.

Firstly, the relativity inherent in the possibility of different space–time partitions prevents this. Consider two observers, A and B. The moving, razor’s edge present of the time-evolution of the universal field is precisely the same for both. Observer A is watching the buzzing fly in our normal partition or scale. He sees the wings as a blur—hundreds of oscillations summed in a single visual display. Observer B, with the process velocity underlying his global dynamics greatly increased, sees the fly barely moving his wings, five wing beats ago being an extremely long time past. The multiple oscillations of the wings (the blur) comprising the ‘present’ for A are in the vastly far past for B. Does B have the right to say that these wing beats are in fact *past* for A, being reconstructed only by the ‘immediate memory’ power of A’s brain? Yet we can imagine an observer C with higher process velocity, in the same position relative to B. All along the time-spectrum of this event we can ‘spread out’ the perception, lowering in a continuous transition the number of oscillations perceived as ‘present’ simply by modulating the ‘energy state’ or process velocity supporting the dynamics of the brain.

Secondly, this analytical approach is the ultimate infinite regress. If one treats motion as a series of states or points (immobilities), one must continuously reintroduce motion between each point to account for the movement. Bergson (1896/1912) argued that this line of analysis ultimately derives from the fundamental partition of the world into ‘objects’ and their ‘motions’ effected by our perception. It is a partition springing from the purely practical need of the body to act—to pick up a ‘stick’ or throw a ‘rock’ or hoist a ‘glass of beer.’ But this purely practical partition is rarified in thought. The separate ‘objects’ become *abstract* space—a network or mesh we place across the concrete extensity of the environment, the meshes of which we can contract at will until they each become a point, and we end with the continuum of spatial points or positions. An object in motion across this continuum is now treated as having a trajectory comprised of a set of these points. Each point or position the object occupies as it moves is now considered to correspond to an ‘instant’ of time, and thus is born *abstract* time—simply another dimension of points in the abstract space. The rarefaction continues. The motions are now treated as *relative*, for we can move the object across the continuum, or the continuum beneath the object. Motion now becomes immobility dependent purely on perspective. All *real*, concrete motion of the universal field is now lost. All quality is lost as well. The motion of the fly becomes a series of the most minute instants of time, each effacing itself instantly before the next, corresponding in fact with the instantaneous death/rebirth of the entire universal field. But on this analysis, there would never exist more than this truly instantaneous 3-D space. Even the brain would have to accomplish its perception *and* its memory storage of the ‘present’ in this same instantaneous slice of time.

But there must be *real* motion. Bergson would insist:

> Though we are free to attribute rest or motion to any material point taken by itself, it is nonetheless true that the aspect of the material universe changes, that the internal configuration of every real system varies, and that here we have no longer the choice between mobility and rest. Movement, whatever its inner nature, becomes an indisputable reality. We may not be able to say what parts of the whole are in motion, motion there is in the whole nonetheless. (1896/1912, p. 255)

He would go on to note:

> Of what object, externally perceived, can it be said that it moves, of what other that it remains motionless? To put such a question is to admit the discontinuity established by common sense between objects independent of each other, having each its individuality, comparable to kinds of persons, is a valid distinction. For on the contrary hypothesis, the question would no longer be how are produced in given parts of matter changes of position, but how is effected in the whole a change of aspect . . . . (1896/1912, p. 259)

The motion of this whole, this ‘kaleidoscope’ as

Bergson called it, cannot be treated as a series of discrete states. Rather, Bergson would argue, this motion is better treated in terms of a melody, the ‘notes’ of which permeate and interpenetrate each other, the current ‘note’ being a reflection of the previous notes of the series, all forming an organic continuity, a ‘‘succession without distinction’’ (Bergson, 1889), a motion which is *indivisible*.

But if this analysis should seem irrelevantly metaphysical, let me remind us of how real it has become for physics. Indeed, if for physics it is true that, ‘‘ . . . a theory of matter is an attempt to find the reality hidden beneath . . . customary images which are entirely relative to our needs . . . ’’ (1896/1912, p. 254), then the abstract concept of space and time described—this ‘projection frame’ for thought—has been the obscuring layer which is slowly being peeled away. Thus the ‘trajectory’ of a moving object no longer exists in quantum mechanics. If one attempts to determine through a series of measurements a precise series of instantaneous positions, one simultaneously renounces all grasp of the object’s state of motion, i.e., Heisenberg’s uncertainty. In essence, as De Broglie (1947/1969) noted, the measurement is attempting to project the motion to a point in our continuum, but in doing so, we have lost the motion. Motion cannot be treated as a series of ‘points,’ i.e., *immobilities*. So Bergson noted, over 40 years before Heisenberg, ‘‘In space, there are only parts of space and at whatever point one considers the moving object, one will obtain only a position’’ (Bergson, 1889, p. 111). And though physics has not attempted to describe positively, as did Bergson, the melodic motion of time, at deepest issue now is a physics embracing *real* motion. Nottale (1996), noting Feynman’s (1965) demonstration that the typical paths of quantum particles are continuous, but non-differentiable, questions the hitherto fundamental assumption of the differentiability of the space–time continuum, applying instead a fractal approach to space–time and its motion, i.e., *indivisible* extents. Bohm (1980), driven by the need to capture real motion (cf. Bohm, 1987), was led to the concept of the holographic field and its ‘holomovement.’ But a property of this field, in global change, i.e., the change of the whole, is an intrinsic nonlocality (cf. Bell, 1987), now experimentally demonstrated. It is a property at odds with relativity. But relativity is in fact the logical end of the classical abstraction. Thus physics continues to work at precisely this abstraction and its peeling away. We must question if cognitive science is immune.

We return then to the 4-D extent of experience determined by natural scaling via the body/brain as a dynamic system. How this extent could exist without a memory storing up discrete states of time was the question. But the question is a question only within the framework of an abstract space and time. The dynamical system that is the brain participates in the real motion of the universal field. This motion is a non-differentiable continuity, best conceived as a melodic flow. Such a flow can support both the qualitative and time-extended aspects of the perceived world, i.e., of experience. There can be ‘buzzing’ flies, ‘rotating’ cubes, ‘stirring’ spoons, and, as we shall see, ‘mellow’ violins.

2.5. Qualitative invariance

Qualitative invariance requires a system whose motion in time is characterized by ‘melodic time.’ Consider the concept of ‘mellow.’ The word has manifold meanings: we can talk of a wine being mellowed with age, a dimension of the word we apply to taste. We speak of a violin being mellow or of a song being mellow, a dimension applying to sound as well as mood. We speak of the interior of a house or room being mellow, referring to the visual. We can say ‘mellow’ of a soil. The concept of ‘mellow’ expresses a very abstract qualitative invariance defined *across* many modalities. At the same time *within* each of these dimensions it is a quality that emerges only over *time*, within the experience of a being dynamically flowing over time. ‘Mellowness’ does not exist in the instantaneous ‘instant.’ This quality can only become experience for a being for whom each ‘state’ is the sum and reflection of the preceding ‘states,’ as a note in a melody is the reflection of all those preceding it, a being whose ‘states’ in fact permeate and interpenetrate one another. If we take this to heart, we should say that the meaning of the word ‘mellow’ is an invariant defined within and across modalities and over time. It is not a homogeneously represented invariant, nor can it exist in space, when space is defined as the abstract, three-dimensional, instantaneous cross-section of time. It means nothing then to store ‘mellow’ as a node in a semantic network, i.e., in a homogenous memory medium, and statically, with relational links to violins, wines, rooms, music, etc. It means just as little to store weights in a connectionist net relating mellow-and-room or mellow-and-violin—again an homogeneous medium with static links, with no support for the time-extent motion required to support such a quality.

We can add a rating test then in the spirit of French:

- ‘‘Rate the song (X) as ‘mellow’’’
- ‘‘Rate Al Gore as ‘mellow’’’
- ‘‘Rate George Bush as ‘mellow’’’
- ‘‘Rate a fly as ‘buzzy’’’
- ‘‘Rate a hummingbird as ‘buzzy’’’
- ‘‘Rate bacon as ‘crunchy’’’
- ‘‘Rate pumpkin pie as ‘crunchy’’’

It may be objected that the computer cannot rate Gore on ‘mellow’ if perhaps it has never experienced Gore. But this was French’s point for all these tests. The required information cannot, *in practice*, approaching only in limit a theoretical impossibility, be pre-programmed; it would have to be experienced over time. But the deeper, correlated question that is riding along here, as we have analyzed what experience is, is whether the computing machine is the sort of device that could ever, *in principle*, experience anything.

3. The broadly computational dynamics supporting experience

In essence, I would presume that the Turing Test is truly envisioned by most as ultimately employing a robot. The computer-as-mind community would rightfully reject the limitation of a blind machine bolted to the computer room floor taking the test. Give the machine its visual input, its auditory sensors, its mechanical arms and legs, i.e., its sensory-motor matrix. Let it walk around behind Mr. Gore for months if it wants to. Let it go to symphonies and listen to violins. Let it try to gather the experience via some yet unknown algorithms. Can the robot/machine, whether harboring connectionist net, symbolic programming or both, and mechanisms for acting tied to the net and/or symbols, yet gain the *experience* necessary to pass these tests? The answer being developed here is ‘no.’ In its fundamental, abstract structure, nothing has changed. It still is a creature of abstract space and abstract time. The perception of multi-modal invariance defined over continuous or melodic time is yet beyond it.

Let me expand the last comments. The computing device, as currently conceived, is a creature of abstract space and abstract time. It is often indeed built on top of a real dynamics—it can employ the real motions of real electrons and use real electromagnetic fields. It is illusory however to be comforted by the thought that a discrete system ‘‘is so only in a fictitious sense,’’ because the dynamics beneath are continuous (e.g., Prinz & Barsalou, 2000). Chalmers (1996) is similar, arguing that a weakness of Searle’s (1980) Chinese room argument is that it does not respect the crucial role of implementation. The program Searle rejects is indeed syntactic, but, Chalmers would argue, implementing the program in the ‘‘right causal dynamics’’ could effectively result in mind and semantics. But, as opposed to the dynamics we contemplated earlier underlying the attractor supporting the brain’s response to the rotating cube, the real dynamics of the computing device has been hitherto in no way structurally related to or time-proportional to the dynamic structure of events of the environmental field surrounding the machine, while the (discrete) operations it is designed to support are independent of any particular dynamics. The operations are syntax, and this syntactic manipulation is supposedly carrying the effective load for thought and perception. We cannot afford schizophrenia here. Either the dynamics is paying the bill, or the syntactic operations are paying the bill. If it is the dynamics, we must understand what it is about the ‘right’ dynamics that pays the bill. For the moment, let us suppose it is the syntax as has largely been supposed and is still at the least deemed to have the critical role as evidenced by Chalmers, or Prinz and Barsalou, etc.

Let us remind ourselves about syntax. Syntax, in its simplest form, can be simply defined as rules for the concatenation and juxtaposition of objects (e.g., Ingerman, 1966). A rule of syntax states a permissible relation among objects. Semantics, in the context of syntax-directed processing, is defined as the relationship between an object and the set of meanings *attributed* to the object, and an object to which a meaning is attached by a rule of semantics is termed a symbol. A syntax-directed processor, it is clear, is *scalelessly* defined with respect to time. Speed up the dynamics, it is irrelevant to the rules for concatenating and juxtaposing objects or the results therefrom. This is equivalent to saying that such a processor can be entirely defined in terms of abstract *space*. This space is prior to 'spaces' such as Euclidean space, Riemannian space, etc. It is a space which relies on a principle of ''infinite divisibility'' as Bergson (1896/1912) put it, and therefore it is also the concept of the infinite continuum of mutually external positions as discussed earlier. As was noted, all *real*, *concrete* motion or evolution of the universe is lost here, and hence all quality.[1]

From automata theory it is known that a device for which the term 'syntax-directed' is appropriate will always be equivalent to a device that applies these syntax rules one at a time since the rules are logically independent, despite what might appear to be simultaneity in their application. Thus the syntax-directed processor can be considered to be moving through a series of *states*. Since the processing can be considered a movement from state to state, it can again be described sufficiently by the concept of abstract or homogeneous time—the line of discrete 'instants.' This 'time' is in reality just the same abstract *space*.

The syntactic device then seems to be in all respects a limiting case—scale-less as opposed to scale, static or spatial as opposed to real motion. The lesser cannot account for the greater. If the device is truly discrete state, it cannot account for motion, falling into the infinite regress earlier noted. It cannot then account for the inherent motion, as indivisible extent, characteristic of consciousness; it cannot span even two states. To argue that it can serve as a sufficient approximation (Prinz & Barsalou, 2000) is to ignore that an 'approximation' may be meaningless. There are phenomena as simple as our buzzing fly for which this device as an approximation is utterly misleading.

But the key here, it will be argued (a variant of Chalmers), is *causally effective* syntax. This is the essence of the 'language of thought' (LOT) framework. Fodor and Pylyshyn (1995) argue:

> If, in principle, syntactic relations can be made to parallel semantic relations, and if, in principle, you can have a mechanism whose operations on formulas are sensitive to their syntax, then it may be possible to construct a *syntactically* driven machine whose state transitions satisfy *semantical* criteria of coherence. (1995, p. 113)

The first, ''If, in principle'' is the problem. Let me paraphrase an example of Fodor and Pylyshyn. From a semantic point of view, according to these authors, ''John stirred the coffee with the spoon,'' as a syntactical structure, would entail a semantic consequence, ''The coffee was stirred by a spoon.'' The semantic, as they argue, is reflective of syntactical structure. But our earlier discussion of the event of coffee stirring has shown that a multi-modal event invariance structure with qualitative invariance is far richer than can ever be fully specified by syntax; a syntactic mechanism cannot be 'sensitive' to this structure. But in the LOT model, not only must it be 'sensitive to,' it must also drive the dynamics of the brain supporting these time-extended invariance structures. Thus the continuously modulated array of energy reflected from the buzzing fly or rotating cube begins arriving at the retina; a 'snapshot' is taken; the symbol manipulation begins, e.g., Marr's (1982) 2.5D sketch, then another snapshot, etc. To be causally effective, we must have a mechanism that is not just sensitive to syntax (i.e., following syntactic instructions), but this syntax would simultaneously have to turn around and drive the dynamics of the

---

[1] Searle, manipulating his symbols in the Chinese room, was criticized for being unrealistic—he would be moving too slow. If he were going at the true speed of a computer, there would arise 'understanding.' It is implied that, mysteriously, faster symbol manipulation brings about a qualitative difference! Somehow, for example, applying the rewrite rule $S \rightarrow NP + VP$ more quickly, results in something beyond $NP + VP$. The only merit in this argument is the underlying intuition that time-scale does matter. Were a normal speaker to say the word ''understand,'' for a listener in a very small scale, the initial ''un'' would vastly precede the final ''nd,'' making language comprehension impossible.

mechanism, e.g., drive a brain-supported attractor underlying a time-scale specific form (cylinder, serrated-edge) of the rotating cube. How all this could happen is a vast gap in the LOT hypothesis. Garson (1998) makes at minimum the case that for driving (or responding to, for that matter) a dynamic, chaotic system, syntax, as causally effective, is severely problematic—the multiple realizability found in chaotic processes undermines any mapping of syntactic causal roles to physical or neural states. But more precisely in the context here, *sampling* (as in strobing) destroys both quality and the time-extended information for form, and we are doing computations here on discrete *samples* or states of the rotating cube. In truth, the syntactic processes know (are reflective of) neither the actual transformation nor the scale of time involved. So how could the syntax ever specify or drive the appropriate attractor? What we are looking for, in truth, is a mechanism for *semantic*-direction. This will include a principle by which the brain is both sensitive to, or driven by, the invariance structure of events and wherein this structure is simultaneously a semantic structure.

Meaning, it was stated, is *attributed* to the symbols of a syntactic system; meaning is extrinsic. The symbols, as Harnad (1990) noted, are not grounded. Yet, simply attaching the symbol strings to actions, as has been suggested (Vera & Simon, 1993), clearly buys us nothing here, in truth being no different than assigning a ‘rule of semantics.’ It does not solve the problem of a quality-less world, it does not solve the problem of a world as experience where *events* are specified by invariants defined only over time-extended transformations, it does not create 4-D extents. Nor does it solve the ‘parasitic’ gain of meaning from an extrinsic source, as Ziemke (1999) points out with respect to so many of these approaches.

The connectionist net does not escape the difficulty, though Sun (2000), rightly hoping to capture Heidegger’s notion of being-in-the-world, argues that a network interacting with the environment qualifies as a direct, unmediated interaction and thus ‘‘the meanings of encodings lie in the intrinsicness of the weights and wiring’’ (p. 165). There are several things than can be said here. I note, though this seems to carry little weight in connectionist circles, that the digital computer and the connectionist net are Turing Machine equivalents (Adams, Aizawa & Fuller, 1992; Chater, 1990), and the syntax-directed device is simply a classical embodiment of the Turing device. Further, associating an input vector with an output vector can be seen as in essence equivalent to forming a syntactic rule, a legal concatenation of objects (vectors). The idea that it is a direct interaction is questionable given that the input vectors are themselves arbitrarily specified, nor does the dynamics of this device reflect or bear a direct relation to the dynamical structure of external events. But this said, let us ask a simple question: how can a set of weights carry the information required to specify the transforming cube? Or the buzzing fly? Or the heron-fly? How will these weights represent a given scale of time? Or a 4-D extent? How, in other words, do weights create the time-extension of experience, without which conscious environment–organism ‘interaction’—inherently time-extended—is a fiction? Let us ask a simpler question: how will these weights, however intrinsic we think them to be, specify the external image or form of the cube? Or the fly? How, even if tied to motor systems? Bickhard (2000) and Bickhard and Richie (1983) would simply call this the fundamental problem of ‘encodingism,’ and ask how these weights (encodings) are unfolded as the perceived world unless one already knows what the world looks like?

The concept that the manipulation of internal representations in some form must carry the weight of semantics has become ingrained. The brain, however, may be as much a dynamic system as an electric motor. The operations we expect a motor to perform are intrinsically a function of its dynamics. It generates a field of force. The ‘mathematics’ of the motor is only a description of the real, concrete dynamics at work. The motor does not exist to implement mathematical operations; it exists to generate a field of force. But this kind of point, revolving around whether computation is a ‘natural kind,’ has already been made. For Revonsuo (1994), ‘‘Brains compute no more than planets solve differential equations about their orbits’’ (p. 259), while to Searle (1994, 1997), syntax (or computation) is no more an *intrinsic* property of the material world than the ‘bathtub-ness’ of some object, but exists only in the eye of the beholder. But the point unfortunately lacks some force of clarity in the absence of a model of the brain that makes it concretely clear. Teng (2000), for example, completely turns Searle's argument around, proposing that we study the 'syntax of the brain,' abstracting the brain's organization into a computational description which then can be implemented on *any* physical system, thus enabling that system to support semantics. But an AC motor will not be instantiated by an abacus, a Searle, or the population of India, i.e., by any arbitrary components. Indeed, the more complex the dynamics, the more constrained the possibilities for instantiation. Ignoring the fact that a computational description of the brain is just what Newell and Simon (1961) thought they had, we would clearly agree that abstracting the 'computational description' of an AC motor for simulation on *any* physical system is meaningless. Tesla would indeed wonder what we were about. But we cannot bring ourselves to think this is meaningless for the brain.[2]

But the critical analysis thus far means little without concretely embedding the thoughts presented on the relation of semantics and the nature of experience within some concept of the form of 'device' necessary to support semantics. I offer an indication in the next several sections as to how the brain could be conceived as computationally dynamic in a broad sense, which, as Copeland (2000) has pointed out, has been left fully open by Turing. Though I shall give no analysis of the nature of this computation, it should be intuitively clear that a broad form of computation is taking place, but that in this, the brain can equally be seen as much a dynamical device as is an electric motor, and that in this dynamics lies a solution to the problem of perception, of symbol grounding, and a basis for semantic direction.

3.1. The external image as virtual action

Let us consider the implications of what we have seen so far in our consideration of experience. On the one side, we had the transforming image of the rotating cube. It is an image defining a temporal *scale* on the universal field. On the other side, we had the transforming neural dynamics of the brain, supporting, as we posited, an attractor. It is a transformation which we know must determine the time scale of the image; it is structurally related; it is even proportionally related. This dynamics, with its proportionality, is a function of the underlying process velocity of the system. Gibson would have termed all this 'resonance.' But we come then to the problem. We see nothing in the brain that can possibly explain the experienced *image* of the cube. We see only attractors, neural patterns transforming. We stand before the famous—the DH-maximized—gap.

Beyond the gap lies Bergson's theory of mind. The dynamically changing, 'kaleidoscopic' field which carries the cube transforming, the fly buzzing, the neural patterns dynamically changing, Bergson (1896/1912) saw as in essence a holographic field (as later would Bohm, 1980). The time-motion of this field is critical. Bergson, we noted, saw that it must be conceived, not as a set of discrete instants or states, but as the motion of a melody, where each 'state' (or 'note') interpenetrates the next, forming a dynamic, organic continuity. Treating the motion of time as a divisible line with 'parts'—'instants,' 'past,' or 'present'—has no meaning in this conception. Time-motion is an indivisible. The 4-D 'extents' of our scales are indivisibles. They do not consist of sets of 'parts' that are born and instantly cease to exist. Within this global motion, 'brain' and 'body' and surrounding 'objects' have no more independent or mutually external reality than the 'particles' of physics. They are abstractions, born of the fundamental partition into 'objects' and 'motions' perception makes in this field. It is a partition meaningful only at a scale of time useful for the body's action.

As did Mach, Bergson saw this field as an immense field of motion or *real* actions. Any given 'object' acts upon all other objects in the field, and is in turn acted upon. It is in fact obliged:

---

[2] I might be accused here of implicitly attacking only a 'narrow mechanism,' ignoring the wider mechanism, as found, for example, in Turing's O-machines (and beyond). Indeed I am focused on the narrow definition that has been prevalent and critiquing its ability to support experience, therefore semantics, but the alternative 'device' I will support fits in a class of wider mechanism, though the properties of this wider mechanism are indeed challenging to define in detail.

> . . . to transmit the whole of what it receives, to oppose every action with an equal and contrary reaction, to be, in short, merely the road by which pass, in every direction the modifications, or what can be termed *real actions* propagated throughout the immensity of the entire universe. (1896/1912, p. 28)

Implicit within this field is an elemental form of awareness/memory. This is due firstly to its holographic property wherein there is a reciprocal response of each field 'element' to every other field element. In a hologram, the information (or wave) from any point of an object is spread across the hologram, but conversely, at any point of the hologram is the information for the entire object. Similarly, the state of each 'element' of the field is reflective of or carrying information for the whole—it is, in a very elementary sense, aware of the whole. Secondly, the indivisible or melodic time-motion of this field is such that each 'state' is the reflection of all previous states. When considered then at the null scale—the most minute possible scale of time—there is already an elementary form of perception defined across the field, in Bergson's terms, an instantaneous or 'pure perception' with (virtually) no admixture of memory. The question now becomes not how perception arises, but how it is limited.[3]

We have tended to take a photographic view of things, Bergson argued (and Gibson would also insist), asking as it were how the brain develops a picture of the external world, or in current terms, how a *representation* is developed and interpreted (or unfolded from a code) as the external world. But he argued in holographic terms:

> But is it not obvious that the photograph, if photograph there be, is already taken, already developed in the very heart of things and at all points in space. No metaphysics, no physics can escape this conclusion. Build up the universe with atoms: Each of them is subject to the action, variable in quantity and quality according to the distance, exerted on it by all material atoms. Bring in Faraday's centers of force: The lines of force emitted in every direction from every center bring to bear upon each the influence of the whole material world. Call up the Leibnizian monads: Each is the mirror of the universe. (1896/1912, p. 31)

Individual perception, he argued, is *virtual action*. An organism is a system of field elements organized for action. Embedded in the vast (holographic) field of real actions, those influences to which its action systems can respond are reflected, as it were, as virtual action, the rest simply pass through.

> Only if when we consider any other given place in the universe we can regard the action of all matter as passing through it without resistance and without loss, and the photograph of the whole as translucent: Here there is wanting behind the plate the black screen on which the image could be shown. Our 'zones of indetermination' [organisms] play in some sort the part of that screen. They add nothing to what is there; they effect merely this: That the real action passes through, the virtual action remains. (1896/1912, pp. 31–32)

Put in holographic terms, the brain is now seen as a modulated reconstructive wave in a holographic field. The re-entrant architecture, the resonant feedback loops, the 'scales' of neural dynamics all ultimately create this modulated wave. As a wave

---

[3] A natural reaction from some quarters will be to label this 'panpsychism,' carrying thus an automatic rejection. But panpsychism justly carries the stigma of merely being a convenient hypothesis or stipulation on the universal field and its objects without much justification. In Bergson's case (or Bohm's), however, attention is being called to real properties of the field, the implications of which must be considered. If the state of every element reflects the whole, if the motion of the whole is indivisible and therefore again, the state of every element reflects the history of the whole field, it is difficult to avoid the concept that an elementary perception is implicit within this field at the null scale. Note also, this is the *null* scale of time. When panpsychism speaks of 'consciousness' in the field, what scale is meant? Does it refer to the scale of consciousness as we know it? Again, in Bergson's case, the whole dynamical apparatus supporting the brain as a wave is required to impose a time-scale on this field in order to support consciousness as we (or even frogs or chipmunks) know it.

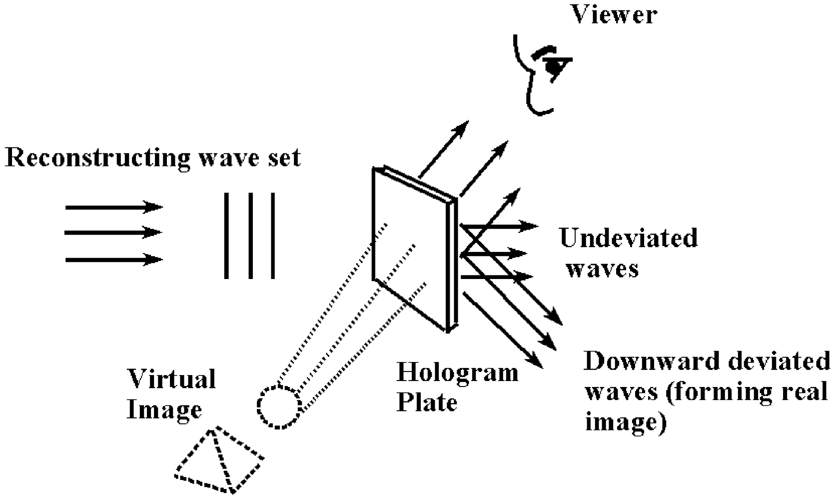

Fig. 4. Holographic reconstruction. A set of plane waves of the same frequency, f_1, as the original reference wave used to store the interference pattern (hologram) strikes the plate and is diffracted in different directions. The upward rising wave set specifies the virtual image of the (stored) objects. Another reconstructive wave modulated to a different frequency, f_2, can reconstruct a different stored wave front, e.g. perhaps the image of a coffee cup, etc. (cf. Kock, 1969).

travelling through a hologram is specific to a virtual image (Fig. 4), this wave is specific to a virtual subset of the field related to the body's possible action. The modulation pattern is determined by the information in the field to which the action systems can respond. This information we have already seen—it is the invariance structure of events described by Gibson.[4]

Conceiving of the brain as a wave is not unprecedented. Globus (1995), discussing the work of Yasue, Jibu and Pribram (1991), describes the evolving brain states ''as best thought of as complex valued wave flows. Constraints on the brain's (state) evolution are elegantly represented by Fourier coefficients of the wave spectrum of this formulation'' (p. 145). Pribram (1991), however, is in general focused on describing the dendritic networks of the brain as mathematical manifolds, and their resonance as 'holoscapes.' As in his original model (Pribram, 1971) of perception, the external image of the 'world-out-there' is somehow 'projected' outwards from recorded wave patterns in the brain, though now (Pribram, 1991) couched in terms of 'projecting invariants' through corticofugal paths. Missing in the analogy is the reconstructive light wave that transduces the recorded (neural) interference patterns into an optical image. Also unexplained is the homuncular eye which now views the projected image. More fundamentally, Pribram yet sees the subject/object relation in terms of space, but this relation is the all-important key.

There is no homunculus in Bergson's model viewing a reconstructed wave front. As he stated, ''*Questions relating to subject and object, to their distinction and their union, must be put in terms of time rather than of space*'' (1896/1912, p. 77). The buzzing fly or rotating cube and the transforming brain are phases of the same dynamically transforming field. At the null scale of time there is no differentiation. But gradually raise the ratio of events in the universal field to events at the highest scale or level of the brain. At the null scale, initially, this E/O ratio would be nearly 1:1, but as it raised, there

---

[4] I focus here only on explaining the source of the 'external' image, located in depth, in volume, in space, and not upon internal perception, e.g. pain, feelings. The source of the former lies in virtual action, but to Bergson, the key to a theory of the latter is that these (bodily experiences) are the field of *real* action.

would gradually form a vaguely outlined ensemble of whirling 'particles,' then the form of the fly would begin to coalesce, then to barely move its wings, then become the heron-like fly, then become the buzzing being of our normal scale. The dynamical state of the brain is specific to a time-scaled subset of the past states of the field, i.e., *it is specific to a time-scaled subset of the elementary perception defined over the entire field*. Symmetrically, because it is a specification of action, the virtual image is simultaneously the display of how the organism can act at this scale.[5]

3.2. The relativity of virtual action

> [Objects] send back, then, to my body, as would a mirror, their eventual influence; they take rank in an order corresponding to the growing or decreasing powers of my body. The objects which surround my body reflect its possible action upon them. (1896/1912, pp. 6–7)

So Bergson would begin his argument that perception is virtual action. The function of the brain is not representation, he held, but the preparation of an array of motor acts. Highly related to Gibson's (1979) notion of the perception of 'affordances,' the perceived world thus becomes the reflection of an array of action possibilities.

The *order* being carved out of the ambient energy flux (Bohm's 'explicate' order) is a particular order defined relative to the action capabilities of the organism. The regularities of the world, the shared commonalities across observers that save us from pure idealism, derive from the invariance laws (in the realist's field) to which these systems can respond. It is worth a reminder here on the large number of findings that have pointed to the general concept that the objects and events of the perceived world are in a real sense mirrors of the biologic action capabilities of the body (cf. for example, Viviani & Stucchi, 1992; Viviani & Mounoud, 1990; Glenberg, 1997). O'Regan and Noë (in press), cf. Robbins (in press) argue in the spirit of virtual action for the basis of vision in 'sensori-motor contingencies,' while Churchland, Ramachandran and Sejnowski (1994) express the importance to visual computation of re-entrant connections from motor areas to visual areas. However, the principle of virtual action may carry an implication deep enough to incorporate—as Weiskrantz (1997) has discussed on the findings of Nakamura and Mishkin (1980, 1982)—the reasons blindness can result simply from severing visual area connections to the motor areas.

As we earlier considered the effects of introducing a catalyst into the dynamical makeup of the body/brain, we already previewed the relativistic aspect of this principle. Let us complete the implications, for the time-scaling of the external image is not a merely subjective phenomenon—it is objective, and has objective consequences realizable in action.

Consider a cat viewing a mouse traveling across the cat's visual field (Fig. 5). We focus first on the

---

[5] Though we can intuitively grasp the notion of a hierarchical dynamics of the brain, and by raising its underlying velocity of process defining a ratio (E/O) relative to the events of the universal field and thus an implicit time scale, an important question becomes how this dynamics is viewed in terms of holography. How is it simultaneously conceived and described as a reconstructive wave which is effecting a time-scaling on the hologram (or holomovement) which it (the brain supported wave) is passing through? I am relying only on the intuition here that the manifold of local resonant fields in the brain (of which there are many if one studies Yasue et al., 1991, for example), at some ultimate level of description, globally comprise a large scale, low frequency wave (relative to micro-events in the field). The higher this large-scale frequency, the smaller the specified scale of time. It would be interesting to see the results if researchers such as Yasue were viewing the brain rather in this Bergsonian version of the holographic metaphor.
But it is not of course this 'simple.' The wave is not just passing through (and constrained) by a diffraction grating or static, recorded interference pattern. Rather it is constrained by the *dynamic* (ongoing, over time) invariance structure of a given event, e.g., the rotating cube, and its relation to action systems. Yasue et al. have indeed attempted to describe how their neural wave equations can support (ongoing) group symmetries. It would be wonderful to see such as these focus initially on just one problem, for example, a rotating cube on a texture gradient, describing how a neural wave equation could incorporate the cube-event's structure of symmetry constraints, how, were it moving forward, size constancy is preserved relative to the gradient, how it could be scale variable—specifying serrated-edged figures as the underlying energy states are lowered—and incorporating the action systems integrally as contributors to this wave.

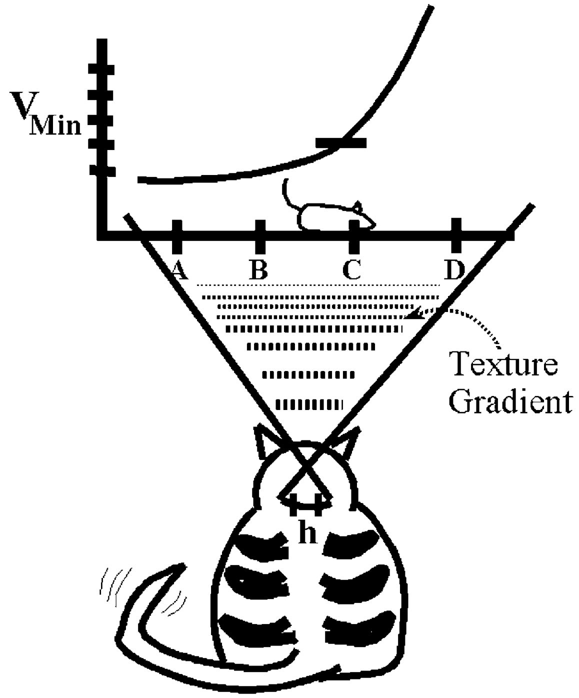

Fig. 5. Hypothetical function describing the minimum velocity required for the cat to intercept the mouse at D. (After Robbins, 2000, 2001.)

Gibsonian structure of this field and its complex projective invariance. There is first of all the texture density gradient stretching from cat to mouse. Were the mouse moving across this gradient towards the cat, the size constancy of the mouse as it moves is being specified, *over time*, by the invariant proportion, $S \propto 1/N$, where S is the (increasing) vertical size of the mouse on the retina, N the (decreasing) number of texture units it occludes ($SN = k$). Were the cat in motion over this field towards the mouse, then 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 value, relating to impending contact with an object or surface, has a critical role in controlling action (Kim, Turvey, & Carello, 1993), and it is implicitly defined then in the brain's 'resonance' state. This entire structure (and much more than described) is supported, over time, by the 'resonant' or dynamical pattern of the brain.

Within this dynamical pattern, there are 'tuning' parameters for the action systems (cf. Turvey, 1977a). Turvey described a 'mass-spring' model of the action systems, where, for example, reaching an arm out for the fly is conceived as in releasing an oscillatory spring with a weight at one end. 'Stiffness' and 'damping' parameters specify the end-point and velocity of such a system. Time is necessarily another parameter. Note that we can translate the mouse and his track towards or away from the cat, and yet the horizontal projection (h) on the retina is the same, any number of such mice/tracks projecting similarly. Therefore, h/t is not enough information to specify unambiguously the mouse's velocity and the needed information required for a leap. The needed muscle-spring parameters must be realized *directly* in the cat's coordinative structures via properties of the optic array, e.g., the texture density gradient across which the mouse moves and the quantity of texture units he occludes.

At our normal scale of time, we can envision a function relating the minimum velocity of leap (V_{min}) required for the cat to leap and intercept the mouse at D as the mouse moves along his path. But how is the velocity of the mouse specified by the body? A physicist requires some time-standard to measure velocity. He could use a single rotation of a nearby rotating disk to define a 'second.' But were someone to surreptitiously double the rotation rate of this disk, the physicist's measures of some object's velocity would be halved, e.g., from 2 ft/s to 1 ft/s. But the body must use an internal reference system—a system equally subject to such changes. This system must be an internal chemical velocity of the body, a velocity it was argued, that can be changed by introducing a catalyst—an operation that can be termed, in shorthand, modulating the body's energy state. If I raise this energy state, the function specifying the value of V_{min} for the cat must change. This is simply to say, with reference to our example, that the perceived velocity of the object (mouse) must be lowered, for its perceived velocity must be a reflection of the new possibility of action at the higher energy state. There is a new (lower) V_{min} defined along every point of the object's trajectory, and therefore the object, *if perception is to display our possibility of action with ecological validity*, must appear to be moving more slowly.[6] If the fly is now flapping its wings slowly, the perception is a specification of the action now available, e.g., in reaching and grasping the fly perhaps by the wing-tip. In the case of the rapidly rotating cylinder with serrated edges (once a cube), if by raising the energy state sufficiently we cause a perception of a *cube* in slow rotation, it is now a new specification of the possibility of *action*, e.g., of how the hand might be modulated to grasp edges and corners rather than a smooth cylinder.

This dynamic system, composed of environment and organism, undifferentiated at the null scale, is truly a tightly coupled, reciprocally causal system. It is a symmetric system, and as Shaw and McIntyre (1974) had pointed out, referencing Mach (1902), such a system is in equilibrium. A change in one half of the system demands a corresponding change in the other to maintain equilibrium. In this case, we have a *cognitive* symmetry, maintaining the equilibrium between the organism's psychological states and the information states of the environment (1974, p. 343). The relativity viewed here is an implication of this symmetry.

3.3. Situatedness and time

There are no representations in this system, i.e., there are no internal symbols within the brain carrying the weight of semantics. The objects of perception, located externally in depth, in volume—the buzzing fly, the transforming cube—are the 'symbols.' These are inherently grounded, for they are reflections of the possibility of action. The system is embedded in time, in the melodic flow of the universal field. Winograd and Flores (1987), in their early argument for situatedness, rejected the view of cognition as symbolic manipulation of representations that are understood as referring to objects and properties in the 'external' world. Following Heidegger's philosophy of being-in-the-world (*Being and Time*, 1927) they noted:

> Heidegger makes a more radical critique, questioning the distinction between a conscious, reflective, knowing 'subject' and a separable 'object.' He sees representations as a derivative phenomenon, which occurs only when there is a breaking down of concernful action. Knowledge lies in the being that situates us in the world, not in reflective representation (pp. 73–74).

Heidegger was certainly aware of Bergson. Cassirer (1957) was straightforward, noting, ''It is the lasting achievement of the Bergsonian metaphysic that it reversed the ontological relation assumed between being and time'' (p. 184). The relationship of subject and object in terms of time constitutes the fundamental framework within which 'situatedness' truly lies. Practically, in terms of constructing a conscious situated robot, it means (at minimum) the following:

(a) The total dynamics of the system must be proportionally related to the events of the universal field such that a time-scale is defined upon this field.
(b) The dynamics of the system must be structurally related to the events of the universal field, i.e.,

---

[6] Why is ecological validity a requirement? Firstly, if it were not, there would be failures of action. We reach out in leisurely fashion, as the barely moving wing of the fly specifies is possible, but the fly is long gone. We are fooled by our perception. Behind this argument is again the concept that as a coherent biological system, a change at any 'level' affects the whole. Such a change is an in principle possibility, and as such, we are asking if nature can have failed to allow for it. Still, I would admit, in specific cases a disconnect between action and perception may be possible, but it would have to do so by affecting some level of organization below which the areas involved in computing virtual action have 'awareness.' As a general rule, it would not be good and nature should have anticipated the variability. The older person driving (cautiously, slowly) down the road, the cars seeming to buzz by at great speed, may not be being fooled by his perception—the high velocity of things happening around him may specify his reduced capacity to act. (The gerontology literature carries this implication, e.g. Birren, 1974; Weiss, 1969; Wallace, 1956.) Secondly, it can be argued that this in principle possibility may in fact already be realized when we consider the spectrum of dynamics across organisms. The energy state underlying the dynamics of the skittering chipmunk or the chameleon who flicks the fly out of the air may well reflect already this in principle possibility. Fischer (1966) gives related considerations across this organismic scale in terms of oxygen consumption per unit body surface relative to life span and metabolic rate.

reflective of the invariance laws defined over the time-extended events of the field.

(c) The information resonant over the dynamical structure (or state) must integrally include relation to/feedback from systems for the preparation of action (to ensure the partition of a subset of field events related to action).

(d) The dynamical structure must globally, in total, support a reconstructive wave.

Nevertheless, there is room in this framework for the representations which Sun (2000) termed 'explicit,' which we know and experience in thought and memory.

3.4. Direct memory and the basis for 'internal' representation

With the dynamics of the brain conceived as supporting a reconstructive wave in a holographic field, Bergson's model becomes the missing Gibsonian model of *direct memory*. But if Gibson's model of direct perception is in effect Bergson's, perception is not solely occurring within the brain. Experience then cannot be exclusively stored there. The body/brain becomes truly the cross-section of a 4-D being in an indivisible time-motion. Bergson (1896/1912) visualized the brain, embedded in 4-D experience, as a form of 'valve' which allowed experiences from the past into consciousness depending on the array of action systems activated. In updated terms, we would say that the brain, embedded in the 4-D holographic field, again acts as a modulated reconstructive wave. Loss of memories—amensias, aphasias, etc.—would be due, as Bergson (1896/1912) in essence argued, to damage that causes inability to assume the complex modulatory patterns required. This does not mean that there is no form of memory stored in the brain. The form of memory Sherry and Schacter (1987) defined as 'System I,' to include the procedural, is obviously brain-based. This includes the sensorimotor 'schemas' of Piaget, where, for example, an object such as a cup becomes embedded as it were in a matrix of possible actions—lifting, drinking, pouring—which are initially overtly acted out when a cup is perceived, but are ultimately inhibited with age. These become a basis for triggering modulatory patterns. The relation between these two forms of memory—that based in the brain and that which is not—is a complex one and a subject for much further theory.

Between perception and memory there is a symmetry however. The same invariance laws which determine the perception of an event also drive remembering. This implies a basic law of the fundamental operation of direct retrieval, the ubiquitous phenomenon of redintegration:

> 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.

The law of 'redintegration' was stated by Christian Wolff, a disciple of Leibniz and a mathematics professor (this explains the term!) in his *Psychologica Empirica* of 1732. In effect, Wolff's law stated that ''when a present perception forms a part of a past perception, the whole past perception tends to reinstate itself.'' Hollingworth (1926, 1928) would devote two works to the subject, and while the term's usage in the literature is sporadic, it is fundamental. Examples of this everyday phenomenon abound. I walk outside and a flash of lightning reminds me of a childhood storm, or a rustling movement in the grass reminds me of an encounter with a snake. 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.

The 'pattern' is the invariance structure of the event E′. This structure moves the body/brain into a modulatory pattern similar to that evoked by a previous event E, such that E will be reconstructed. Events, we have seen, are defined by a structure of transformations and structural invariants. The more unique this structure, the easier it is to reconstruct the event. It is exactly as if a series of wave fronts, w_i, were recorded upon a hologram, each with a unique frequency (f_i) of reference wave (Fig. 4). Each wave front (or image) can then be reconstructed uniquely by modulating the reconstructive wave to each differing frequency, f_i.[7]

Thus suppose a series of perceived events, e.g., a man stirring coffee, a baseball hurtling by one's head, a boot crushing a can. Each has a unique invariance structure. To create the reconstructive wave for these, i.e., to evoke over the brain the needed modulatory/dynamical pattern, I might use as a 'cue' respectively—a stirring spoon, an abstract rendering of an approaching object capturing the composite tau value (Craig & Bootsma, 2000) of the original baseball event, and an abstract rendering of one form descending upon and obscuring another. But these events are multi-modal and the four-dimensional extent of experience is multi-modal. There are auditory invariants as well defined over the events. Our cues could become respectively—the swishing or clinking sound of stirring, the 'whoosh' of the passing baseball with frequency values capturing its inherent doppler effect, the crinkling sound of collapse of a tin structure. And in the dynamics of the haptic component of the event, we could cue our stirring event by wielding a 'tensor object' that captures the inertia tensor (invariant), I_{ij}, specific to spoon-stirring (cf. Turvey & Carello, 1995).

One can imagine then a quite fearsome paired-associate paradigm as far as verbal learning experiments are concerned. A list would look as follows:

SPOON–COFFEE
SPOON–BATTER
SPOON–OATMEAL
SPOON–BUTTER
SPOON–CORNFLAKES
SPOON–PEASOUP
SPOON–CATAPULT
SPOON–CHEESE
SPOON–TEETER TOTTER
And so on . . .

It is fearsome from a verbal learning perspective since the same stimulus (cue) word appears constantly, thus providing absolutely no clue to which response word is intended. It is the extreme case of the A–B, A–C list paradigm where two lists of words are learned successively, where the pairs share the same stimulus word, for example:

List 1 (A–B)	List 2 (A–C)
SPOON–CUP	SPOON–PLATE
BOAT–LAMP	BOAT–TABLE
KNIFE–SOAP	KNIFE–MEAT
And so on . . .	

---

[7] Though appropriate earlier in the perception portion of this discussion, this note requires the notion of redintegration. Perception clearly is subject to illusions, i.e., misrepresentations, e.g., the Poggendorff, the Ponzo, etc. Gibson (1966, 1979) long argued that these are artifacts, that given an ecological environment, rich with information (invariants), these do not happen, the experience being again 'directly specified.' (In fact, some 'illusions' would count as quite valid percepts from a Gibsonian perspective.) Irvin Rock (1984) opposed Gibson at every turn, trying to show that inferences or mental operations are involved. Typically his experiments involved information-deprived experimental setups, destroying Gibson's texture gradients for example by forcing the observer to judge distances (say of two rods located at different distances on the floor) when looking into a darkened room through a peephole. He would then show that inferences or mental operations must be involved. Norman (in press) has recently argued in detail that the two different conceptions appear to correspond to the dorsal processing stream (Gibson) and the ventral stream (Rock). Both are in communication, but the ventral appears to be engaged critically when information is insufficient for the dorsal, leading to the dominance of 'judgements' (where knowledge of the world is brought in) in the environmentally deprived case.

Beyond this, there are misrepresentations where we see a 'bear' at night that turns out, on moving two steps closer, to be a tree stump. Here Bergson comes to the fore. Note firstly that in stating that the brain's dynamical state is specific to past states of the field, we are implicitly stating that perception is already, simultaneously, a memory (Edelman's 'remembered present,' 1989). Secondly, the redintegrative or direct recall model also implies that the invariance structure of a 'present' event is simultaneously creating a wave that is reconstructive of past events with similar structure. Bergson (1896/1912) thus argued that perception is always permeated with memory experience (he saw the flow of memories to perception as a 'circuit'). The initially indistinct words to a song, with a bit of a clue, from then on are perceived as 'perfectly clear' whenever we subsequently hear the song. The fat form of the tree stump, a bit indistinct at night, is enough to redintegrate a 'bear.' The ecological value of this setup is clear: Old Scarmouth, the local 12.5 lb. largemouth bass, sees the glint of the hook and a certain pattern of dangle of the worm and sees it instantly in context of the redintegrated experiences of his defeats of unworthy fishermen.

Here, theory (e.g., Marschark, Richman, Yuille & Hunt, 1987) focuses on 'inter-item' relations as critical to help us, e.g., we might notice that List 2 is mostly about eating-related response words. The ecological case is far simpler and it is primary. Assume that the subject concretely acts out each event of the 'fearsome' list—stirring the coffee, stirring the batter, scooping/lifting the oatmeal, the cornflakes, cutting the cheese. To effectively cue the remembering, the dynamics of each cue-event must be unique. An invariance structure in effect implies a structure of *constraints*, correlating strongly with the Constraint Attunement Hypothesis of Vicente and Wang (1998) and Vicente (2000). The constraints of the cue-event may be parametrically varied, where increasing fidelity to the original structure of constraints of a given event corresponds to a finer tuning of the reconstructive wave. The (for example, blindfolded) subject may wield a tensor-object in a circular motion within a liquid. The resistance may be appropriate to a thin liquid such as coffee or to a thicker medium such as the batter. The circular motion may be appropriate to the spatial constraint defined by a cup or to the larger amplitude allowed by a bowl. We can predict that with sufficiently precise transformations and constraints on the motion of the spoon (either visual, or auditory or kinesthetic or combined), the entire list can be reconstructed, i.e., each event and associated response word. Each appropriately constrained cue-event corresponds to a precise modulation (or constraint) of the reconstructive wave defined over the brain.

It should be fairly apparent that representing the paired-associate learning experiment as a process of associating two vectors, e.g., one for SPOON (input), one for BATTER ('output' or 'response'), is not close to supporting the actual dynamics that is the case for any one of these events, e.g., the spoon stirring the batter. The former concept of things can be termed a syntactic association as it is in essence an arbitrary concatenation rule. The latter (the dynamic, multi-modal event) is a semantic association where the words (spoon, batter) only very partially indicate or describe the dynamical laws of the event in which spoon and batter are natural participants, and where the appropriately constrained dynamics of the cue-event is required for redintegration.

The SPOON–CATAPULT and SPOON–TEETER TOTTER pairs in our list are in essence analogical events (cf. Dietrich, 2000). Concretely, the subject may have used the spoon to launch a pea at the experimenter, or the spoon may have been balanced on the edge of a bowl. Note that this event is equivalent to a French rating game, all of which were in fact exercises in analogy:

- Rate *spoons* as *catapults*.

The structural invariants of the spoon which support the 'catapulting' invariance structure emerge under this transformation. Harkening back to an earlier example such as KNIFE–SPOON, where the subject used the knife to stir a cup of coffee, then we have an analogical event equivalent to the rating game:

- Rate *knives* as *spoons*.

Placing the appropriate dynamism/constraints on the knife to create a cue event will reconstruct the stirring event. Fundamentally, this redintegrative or direct recall mechanism lies at the basis of analogical reminding, and the rating games of French as well. Analogy itself can be viewed as a form of transformation under which features emerge.[8]

---

[8] Dietrich would consider the work of Gentner (1983), and Falkenheimer, Forbus and Gentner (1989) as a possible mechanism behind this retrieval/reminding. In this approach, the Structure Mapping Engine (SME) treats analogy as a mapping of structural relations. 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 et al. (1992), Mitchell and Hofstadter (1995), and Hoftstadter (1995) level a heavy critique upon this approach, noting the helplessness of SME without this precise setup of features and relations beforehand, and with this setup given, the purely syntactical, nearly 'can't miss' algorithmic procedure that follows. The resultant discovery of analogy is, to quote these critics, a 'hollow victory.' Dietrich himself concludes that a system that supports the emergence of features *at the time of the analogy* is necessary, rather than mapping two concepts with pre-existing features (though he holds SME capable of modification for this).

3.5. Abstraction and redintegration

The modulatory pattern defined over the brain and supporting these invariance structures can be conceived as a continuously modulated reconstructive wave traversing 4-D extended and multi-modal experience. Recall Gelernter's (1994) operation of taking a 'stack' of events across which the invariants stand out. One may conceive of the basis for a 'concept' as a wave of less than perfect coherence supported by the dynamics of the brain (e.g., a composite of f_1 and f_2 in Fig. 4) reconstructing a composite of images or wave fronts (stirring-events) across 4-D memory, over which the invariants across the images/events stand out. 'Stirring' itself, as a *concept*, is an invariant across multiple stirring events in 4-D memory as defined by this operation. In this sense, the operation of redintegration or direct recall is the basis of abstraction and in turn of compositionality. *Language* would then become a mediating device for moving the brain into these dynamical patterns. The sentence, ''The man stirred the coffee'' can be seen as a device to move the brain into a dynamic modulatory pattern supporting the multi-modal invariance structure defining this 'coffee-stirring' event.

This concept of abstraction, where abstraction is achieved by activating a large number of similar events or memory 'traces,' is often termed *exemplar* theory (cf. Crowder, 1993). It has gained significant support in memory theory. Semon (1909/1923) espoused it, having noticed it in Galton's (1883) example of an abstract face derived over multiple superimposed photographs. Goldinger (1998), noting Semon, uses Hintzman's MINERVA 2 (1986) to defend the thesis in detail. Smolensky (1995) in essence comes at least close to this as well, speaking of a 'coffee cup' as a family resemblance across activation patterns. Harnad's (1990) categorical invariants derived from sensori-motor interaction may be viewed as implying this. But note that this requires a device that can store the totality of experienced events in all modalities, in *complete* detail. The theory sees all spoken exemplars of a given word, e.g., of 'spoon,' as stored, and upon re-presentation of 'spoon,' all similar or partially redundant traces being reactivated to create the context free abstraction, SPOON. But as Goldinger (1998) shows experimentally, one does not and cannot simply store only this precise exemplar, but rather *detailed* episodes, to include all the accompanying particulars of voice, inflection, pronunciation, tone. But words are arbitrary phases (identified by us) in extended acoustic wave forms, i.e., whole sentences, including again prosody, inflection, tone, etc. So the whole sentence must be stored in all its aspects. But the sentences are acoustic forms inseparable from the complete, extended multi-modal events of their context, e.g., while eating at the kitchen table or canoeing down the lake. There appears to be no way to halt short of 'storing' everything.

This 'storage halting' problem was implicit in our 'rotating cube' event where it was asked what sampling rate or sample set could be taken and yet preserve the time-extended information defining the event. The same is true of the dynamic transformations defining the coffee stirring event. Preserving the qualitative aspect of these events presents the same dilemma. How is 'mellowness' preserved across anything less than the entire event, i.e., how by using a few samples? Or more simply, a series of 10 notes is played with a constant interval between each note, defining a certain quality. The same series is again played, but one note is held slightly longer, defining a slightly different quality. The qualitative difference is instantly noticeable. The whole of each series must be 'stored' to preserve this quality.

This operation, comparing whole events with their structure, to detect change, is problematic for any memory storing less than the whole of experience. This storage problem is what exemplar theory's opponents, namely the 'abstractionists' (cf. Goldinger, 1998), would avoid. Barsalou would store a 'biting' transformation as three schematic states—''a mouth closed next to the object, followed by a mouth open, and then the mouth around the object'' (Barsalou, 1993, p. 53). But suppose three 'states' were stored of a rotating cube event. Then assume on a subsequent event, the cube is bulging in and out. Recall that the form of the cube was specified by the symmetry information preserved across the transformation, and sampling (strobing) destroyed this information. For a standard observer, comparing against the two wholes, the difference between the two events is immediate. But the discrete sample method observer could in principle sample three states of the 'bulging' cube and in fact 'match' his three stored states of the previous, normal rotating cube, detecting no difference. The 'stored states' method begs the description of change. Similarly, suppose we have observed a series of n coffee stirring events with a similar invariance structure. In event $n+1$ the (normally constant) cup is now bulging in and out. The dynamic invariance structure defines a similar enough modulated wave to reconstruct the first n events, but an 'interference' is automatically specified relative to the cup detail.[9] Redintegration here carries a complimentary form of 'discrimination.' We do not need 'frame axioms' here (cf. Morgenstern, 1996) to tell us the cup should have remained the same, or that 'snap, crackle, pop' is not a legitimate acoustical accompaniment, or that a melting spoon is anomalous, etc.[10] Indeed we would need frame axioms for every detail in every modality, and change statements that could (truly) specify the transformation dynamics (at every scale of time) as well, e.g., what the 'stirring' flow field should look like, or the 'wielding' of the spoon feel like. Every change of detail of the previous n similar stirring events, 'static' and dynamic, is subject to detection by this 'interferometric' method.[11] Thus Jenkins, Wald and Pittenger (1978) describe how, while observing a series of slides of a room, they detected a slight, nearly indefinable difference in one slide. The difference ultimately proved to be a detail—a shadow angle in that slide. But this capability is only available if we have a memory capable of storing whole events in full detail.

We are led again toward the multi-modal, 4-D model of memory being described here where the global flow of the holographic universal field is indivisible, where perception/experience is not occurring solely within the brain, and therefore not being solely stored there. Therefore, no strain exists on 'storage capacity.' And now abstraction is effected by the brain's dynamic modulatory pattern defining a reconstructive wave through this memory.

Understanding the sentence, ''The man stirred the coffee,'' is not a matter of concatenating a set of these abstract, 'compositional' concepts—man, stirring, coffee—via syntax. The linguistic string, with its syntax, is a mediating device causing a modulatory pattern specific to an invariance structure defined across 4-D experience. The linguistic string expresses little of the immense richness of the multi-modal experience with its invariance running throughout. A vast amount of this structure is unspecified—the instrument involved, the force applied, the sound occurring, the aroma, the motion of the coffee, its properties as a liquid, and more. Yet it can be argued that the concept of 'stirring' referenced by the language incorporates all of this. In so far as

---

[9] The analogy here is to the interferometric property of holography. For example, a hologram of a tire can be made at t_1, the tire then subjected to stress, then a hologram made at t_2, and superimposed on the first. An interference fringe would indicate a defect in the tire.

[10] The 'frame problem' (McCarthy & Hayes, 1969) was defined in the context of the situation calculus and fundamentally deals with the problem of change as a result of an agent's action. If the agent puts a red block, Block A, on top of Block B, the axiom [Holds(s,Red(b1))⇒Holds(Result(Puton(b1,b2),s),Red(b1))] allows the inference that the color remains invariant. But a vast number of such axioms need to be stated—the president remains the same, the house is on its foundations, the sun is yet in the sky, etc., all of which must be checked. Imagine what would be required even in the limited frame of the coffee-stirring event for specifying what changes, what remains the same—in every modality. This need for vast numbers of axioms to prove that most things remain the same as actions are performed is the frame problem. (The predictive problem—how things change—is complementary.) A solution to the frame problem is generally considered precisely this: a method to reduce or remove the frame axioms. From the perspective here, the problem showcases an interesting ('interferometric') feature of perception/memory, but is itself a pure artifact of attempting to treat dynamic, multi-modal events symbolically.

[11] A statement in the situation calculus, e.g., [Holds(Result(Stir(Coffee),S0,Mixed(Cream,Coffee)))], specifying an initial situation and the final result of an action, has not yet begun to do the job. Each 'stage' along the motion would need to be specified (again, at what scale?) and for every modality. (How the fly looks, moving from point A to point B is different, depending on scale, or the form of the cube as it rotates.) How will the coffee's flow field transformation be specified? The acoustical quality? This is the 'exemplar' of the abstractionist, stored-state approach. Now each state of the motion, as in Gray's (1995) hippocampal 'comparator,' must be checked against each predictive statement as the action progresses (not forgetting all the frame axioms) to make sure things are going as expected. The sampling (strobing) problem with the rotating cube already speaks against this, but if nothing more, the sheer overwhelming weight of the discrete representation of these dynamic transformations makes this an impossibility.

'stirring' is an abstraction or invariant form across multiple experiences, it gains an individuality or compositional form, but it is far from the abstract, end-element of a Chomsky sentence diagram, and its compositional basis is quite different from what Fodor originally visualized (though quite in tune with his latest (Fodor, 1998) intuitions, where ''Having a concept is something like 'resonating' to the property that concept expresses''). The abstraction 'stirring' yet rests entirely upon, or is defined across the experiences comprising 4-D memory, and without this basis, has no meaning. Thus, ''The man paddled the canoe across the lake,'' with its components and (dynamic) syntax can be viewed somewhat as a musical score (cf. Verbrugge, 1977) used to create a modulatory wave specific to an invariance structure defined across the vastly rich, 4-D, multimodal experience. It can be visualized again as a wave cutting through this experience. But this 4-D experience of canoeing is a seamless whole. It does not consist of parts. This is why any of the multitude of possible invariants that are violated can cause the feeling of anomaly, e.g., ''The water hissed and sizzled as he paddled.''

Though I have argued that this form of abstraction provides a basis for compositional elements, compositionality and systematicity go hand in hand. How do we learn to use these elements in structured patterns? Like Petitot's (1995) dynamic syntax, the systematic rules for composition also seem to be carved out of dynamic flows, which is to say that these too are invariance laws. For Petitot, it can be said that a standard syntactical form, such as xRy, is an invariance law of a high order, where x and y are in some dynamical relation, R, e.g., John stirs the coffee, Joe catches the ball, the bear enters the cave. (This form, as noted, must serve as the 'score' or driver for modulating the larger invariance structure defining the event.)

This dynamical approach to compositionality was the essence of Piaget's approach. Consider his simple experiment on children aged 3–7 (*The Child's Conception of Movement and Speed*; Piaget, 1946). Here three beads are strung on a wire which can be fitted into a small cylindrical 'tunnel.' The beads are of different colors, but we will call them A, B, and C. The beads are run into the tunnel and the tunnel semi-rotated from 1 to N times. A series of questions are asked, ranging from a simple, ''What order will they come out?'' after one semi-rotation, to the ultimate question on their order after any (n) number of half-turns. The child comes to a point of development where he can imagine the consequences of a 180° rotation which moves ABC to CBA and another 180° rotation which moves things back again to ABC, i.e., an invariance of order under a 360° rotation. When now asked in which order would the beads come out when the tunnel is semi-rotated five (or four, or six, or seven, etc.) times, he evidences great difficulty. Some children appear to be exhausted after imagining three or possibly four semi-rotations, and they become lost when jumps are made from one number to another. As Piaget notes:

> . . . But since the child, upon each half turn, endeavors to follow the inversion in every detail in his thoughts, he only gradually manages accurately to forecast the result of three, four, five half turns. Once this game of visualizing the objects in alternation is set in train, he finally discovers . . . that upon each half-turn the order changes once more. Only the fact that up to this upper limit the subject continues to rely on visualizing intuitively and therefore needs to image one by one the half-turn, is proved because he is lost when a jump is made from one number of half-turns to any other. (1946, p. 30)

After this gradual perception of a higher order invariant (the 'oscillation of order') defined over events of semi-rotations, there comes a point then when the child can easily answer the ultimate question for the resultant order for any n-turns. Piaget's explanation, describing the 'operational' character of thought, is foundational to his theory and its 'group' operations:

> Operations, one might say, are nothing other than articulated intuitions rendered adaptable and completely reversible since they are emptied of their visual content and survive as pure intention . . . In other words, *operations come into being in their pure state when there is sufficient schematization.* Thus, instead of demanding actual representation, each inversion will be conceived as a potential

> representation, like the outline for an experiment to be performed, but which is not useful to follow to the letter, even in the form of performing it mentally. (1946, p. 30, emphasis added.)

Thus according to Piaget, operations, freed of their imaginable content, become infinitely compositional. This becomes the basis for forecasting the result of *n*-turns, and it takes the child to about the age of seven. The operations become the generalization of actions performed through mental experiment. This is not simply abstract rules and symbols. As we have seen, these 'schematic' operations are built upon and do not exist without the dynamic figural transformations over which invariance emerges. They are the result of a dynamical developmental trajectory incorporating these figural transformations which requires on average seven years.

I hope I can be forgiven then when I say that the theory of this form of dynamically embedded compositionality and systematicity has a long way to go, but at least the 'device' being described here provides a beginning basis for the dynamical imagery supporting invariance involved and its 'modulation.' It is in the context of this form of a device, I believe, that Piaget (or Petitot) and his compositionality must be understood.

Finally here, let us note that when the wave supported by the brain is functioning as a reconstructive wave, it is acting to re-establish, or is specific to, an original environment–organism field relation of the 4-D holographic field. Note that an essential symmetry assumption has been made implying a very specific dynamical structure supported over the brain and its action tuning parameters which is reflective of the invariance structure of an event. This structure will provide constraints on the characteristics of this wave when described at the neurodynamical level, or the quantum level, or whatever level of the brain's hierarchical scales one chooses. We should view the *global* dynamics of the brain as comprising this wave. We do not see retrieval processes fetching stored elements—object 'features' or 'schematized' objects or events—and re-assembling them as an 'image' or experience, viewed somehow by an homunculus in the brain. Nor do we imagine waves coursing through the brain, reconstructing images/wave fronts *within* the brain, or re-projecting images/wave fronts outside the brain, again for an homunculus to view. Body/brain and 4-D universal field comprise a coherent system. The changing dynamical pattern of the brain modulates virtual objects in time. If the modulatory pattern is sufficiently precise, these may be experienced as images (for example, ''a knife cutting a tabletop''), or depending on the order of invariance (level of abstraction), may be increasingly image-less (as in, ''the utensil interacting with the furniture''). The debated representational status of the brain's dynamical patterns—the attractors, bifurcations, etc.—supporting these invariance structures is given clear place in this model. If we must still call them 'representations' (and I would not), they are clearly in the relation of the part to the whole. They cannot be equated with the whole of thought. Thought is comprised of the simultaneous relation of dynamical patterns with virtual objects of the four-dimensional mind.

3.6. Voluntary action

The canoeist, paddling down the lake, as earlier described, is largely an adaptive system. The representation required for guiding the canoe in the required line is in fact laid out in the presentation, i.e., via the perceptual field—auditory, visual, haptic. Before him lies the lake surface and optical flow field, all 'directly specified' (in the Bergsonian spin this term has now been given) and inherently 'semantic.' It is the field of action, and the paddler is the essence of Heidegger's 'concernful action' or pre-representative thought. It is the perception/action cycle explored by Hurley (1998), now with its deeper dimension of virtual action. The velocity of the waves moving towards the paddler reflects his scale of time and ability to act. But this multi-modal field is not represented (solely) in the brain.

It was this tightly coupled system, E↔A, that seemed most amenable to the treatment of the DH. The emulative representation, A↔E′, posited to handle planning and representative or imaginative thought, has now been seen amenable as well to the DH when treated in Bergson's and Gibson's larger frame. The dynamical patterns of the brain support the modulation and reconstruction of virtual objects of 4-D extended memory. Via a 4-D memory and its redintegrative mechanism, there is at least the basis for the bodily image of (future) actions, the image of E (E′) or the future transformation of the environment as the result of an action, and goals insofar as the goal is represented via an image. The images of past events and actions can at least now find a basis for 'storage' and reconstruction for later use in action. It is here that we meet the ideo-motor theory of action. Expounded of course by James (1890), even Bergson (1902/1920), it has recently been extensively defended by Jeannerod (1994). The image, be it visual or kinesthetic or both, is seen as holding the plan for the action. If I want to learn to dance a certain step, the visual image provides the schematic outline, while the kinesthetic images of the components already familiar to us—walking, turning on one's toes, lifting arms—provide the elements that must be initially integrated. With practice, the integrated and articulated kinesthetic or motor image provides the plan for a fluid act.

But throughout Jeannerod's more modern account, the image is problematic. Firstly, it is assumed that the image—visual or motor—is generated by the brain. How this could occur, and how the homuncular regress could be avoided is unknown. If it is stored in component form and generated by the brain, how is it more than epiphenomenal? Why is the image needed at all, in a causal role, to guide action? Why is it not redundant? This difficulty is an echo of the reduction of the image to elements within a data structure (e.g., Pylyshyn, 1973; Kosslyn & Koenig, 1992), and the claim that by regenerating (with appropriate structuring) these elements, images could be thus accounted for. But we ask now how this *encoded* information is unfolded to the phenomenal experience of the image? Again, Bickhard's (2000) problem of encodingism. And we ask who now views these images? The same regress that plagued the perceptual image now begins for the mental image.

Bergson's theory, conjoined with Gibson's, can explain the reconstruction of the image of a past experience via an event occurring in the environment; it can at least logically support the phenomenal experience of the memory image. This at least underlies the 'content of intention,' for as Jeannerod notes:

> Motor imagery would represent the result of conscious access to the content of the intention, and the content of the intention would constrain the expression of the image. (1994, p. 190)

The content of the intent, e.g., to place your cup on the airline steward's tray, can indeed be abstract, but as already noted, all abstraction rests upon concrete experiences or events, just as in Piaget's operations.[12] Actualized, under the right conditions, as Jeannerod shows (pp. 190–191), and similar to that noted above on the concreteness of images, the intent can be experienced as the motor image.

But automatic event reconstruction is a passive aspect of consciousness or conscious memory. The problem of voluntary action faces the other direction. It is the dynamic, positive aspect of consciousness. The first of the deep questions of voluntary action is this: How is the image (as the content of the intent) summoned or projected voluntarily or 'at will' within the 4-D extent of being and into the body's field of action? Already we have the question of the direction of causal effect. Is the image causing the physical actions, or are the physical actions causing the image? For the dreaming cat, flicking his tail and twitching his paws, are the dream images of the mouse driving the physical effects, or are the dream images merely epiphenomenal? Though Jeannerod would try to have it both ways, for Bergson, the direction of effect was from the image. This is a greater 'hard' problem and I do not claim to answer it here. The course that would need to be followed, I believe, demands a far deeper understanding of the body/brain and 4-D memory, in total, as a coherent system, and of the global motion of the universal field in which this system is embedded. Here physics and psychology meet. Thus to move a step in the speculative vein, indicating perhaps the wider context of research that might be needed to get a grasp of this problem, it is interesting to note the literature on the real effects of mental imagery practice on motor skill (see Jeannerod, 1994, for a review). Then

---

[12] Cassirer (1957, chapter 6) argued extensively that failures of intent, e.g. in disorders such as apraxias, come from a failure of abstraction (or of the symbolic function) where group operations (exactly similar to Piaget's) are no longer supported due to brain damage. In fact, it can be argued that Cassirer was describing, in the sphere of action, the degenerative inverse of Piaget's development of operations.

we can note the effect of this same form of imagery practice in hypnotically altered temporal states (Cooper & Tuthill, 1952; Cooper & Erickson, 1954) but where there now appears to be objective effects on the velocity of action. Where earlier we discussed changes which should exist in the context of changing the space–time partition via a catalyst on the E↔A side of the schema, this latter is now from the A↔E′ side and purely via mental imagery. Whether this would be a fruitful line of research remains to be seen, but it carries a certain symmetry..

Correlated with this problem, let me note again that the emulation model required a control process, C, determining the attention applied to either of the sides of the E↔A↔E′ schema, i.e., an ability to de-couple, so to speak, A from E and attend to E′. Bergson (1896/1912) was quite clear in the context of A↔E′ or imagination (also hypnosis) that the control mechanism required *suppression* of the ‘call to action’ inherent in E↔A, i.e., abstraction from the inherent ‘call’ of the motor state/virtual action intrinsic to a present perception (as virtual action). Glenberg (1997) has also emphasized this suppression mechanism in the context of imaginative memory. Something must break the tie to the present and allow ‘attention’ to focus on, for example, the modulated images of thought. Though we are far from any theory of C, it is equally a natural and a necessity in the Bergson/Gibson framework.

4. Summary of the semantic-directed processor

I will try here to summarize, though not in any way formally, what might be meant by the term ‘semantic-directed processor,’ given what has been described above. The first thing that must be stated is that, in describing this device, the system comprising environment (E) and organism (O), or the E–O field, is the reality. Secondly, this field develops in real or melodic time, as opposed to abstract time. It is only this form of flow that supports the qualitative aspect of the experienced world. Semantic-direction is obtained from the reciprocal interaction of the invariance laws defined over this field-flow and systems for effecting action. Thus an *event* is defined as a local transformation of the field over some limited time. The *syntax* of an event is defined as the set of transformations and invariants defining the event and rendering E an incipient or virtual action for O. *Semantics* is defined as the symmetric relation of E to O. Meaning is a function of this relation, and thus a *rule of semantics* defines a permissible relation between E and O. An E–O event in which a meaning is defined is also an *affordance*. A *semantic-directed processor* is an E–O field wherein the symmetry relationship is *scaled*, such that the ‘syntax’ of O is symmetric to the ‘syntax’ of E with respect to a definite scale of time.

In the syntactic device, since we are dealing with the manipulation of abstract symbols in an abstract space, only a single homogeneous medium of representation is required for these objects and manipulations. In practical fact, all present day computers represent their data in a single homogeneous medium. Though the device may in fact receive input from several media, for example a camera, a microphone, a pressure sensor, a thermal sensor, yet the information carried over each energy form is transduced and represented as a set of abstract objects in a homogeneous medium. (This is true for networks as well.) Symbolic operations then are conceived within the syntactic framework as the manipulation of objects in a homogeneous space. This intuition of a homogeneous medium underlying the representation of events, including the multimodal ‘World-Out-There,’ underlies the equation of the brain with the computing machine as syntactic devices.

Since in the semantic device, the E–O field is the reality, there is no need to make the medium to which ‘input’ is transduced bear the full burden of representation. Events take place within a *medium* or set of media of the environmental field. This may include the optical, sonic, muscular, thermal, and more. The structure of events within these media is given by systems of transformations and invariants defined within and across these media and existing only over time, i.e., systems which are time-extended. The semantic device’s ‘representation’ is multimodal. Since the motion of this field must be treated as continuous and indivisible, the representation, which in fact *is* the E–O field, is four-dimensional. Thus far then, the semantic-directed processor must be characterized in terms of real time and extensity, as opposed to abstract time and abstract space.

It must be characterized in terms of *quality* as well, a phenomenon that arises (at least in part) from the time-scale definition effected by this processor. We have seen earlier, for example in the discussion of mellow, that quality can only be defined over melodic time. The quality of an event, e.g., the mellowness of a sound, is time-scale specific. As we change the time-scale, so the experienced quality will change. It will be 'mellow,' but changed nevertheless. At the null scale we meet the limiting case of an essentially quality-less world, for at this point we deal with the instantaneous states of the universe. We can thus imagine each instant coming into being, instantaneously effacing itself before the next, leaving us with never more than a *single* pulse. But if quality demands the dynamic addition of instants, where each instant interpenetrates the next, each being an expression of those preceding, then the null scale corresponds to *quantity*. The concept of quantity, Bergson (1889) noted, demands just that we strip a multiplicity of objects of that which makes them different, so that we can treat them alike for the operation of counting.

The notion of the quality-less continuum is again the notion of abstract space. Thus the syntax-directed processor, being defined in terms of space and thus quantity, is simultaneously unable to deal with quality. Quality is the prerogative of the semantic-directed processor. As we place succeeding scales on the universe as visualized at the null scale, the qualities are transformed relative to each previous scale of time as each sums up a greater history. The syntax-directed device becomes a limiting case of the semantic device taken at the null scale of time.

4.1. The syntactic problem—the origin of the object language

As Narasimhan (1969) pointed out early in the game, the standard approach of computer modeling is to define two languages, an *object language*, describing the problem environment, and a *meta-language*, used to construct a program in the object language with specific input/output properties. What makes this approach non-trivial, he noted, is the intrinsic complexity of the problem environment which precludes the possibility of constructing uniform methods for solving all problems specifiable within a given environment. The meta-language embodies a meta-framework of primitive solution methods, rules for evaluating solutions, and for constructing to a limited extent further procedures. But to Narasimhan, the real problem was defining the object language in the first place, i.e., creating semantic theories of the world.

The object language/meta-language definition has been the approach for 30 years, always in variants (precisely because the problem environments chosen do differ) from GPS (Newell & Simon, 1972), Minsky's (1975) frames, to Schank's (1975) schemas and beyond. French (1999) levels a heavy critique of recent variants of this approach, e.g., BACON (Langley, Simon, Bradshaw & Zytkow, 1987), SIAM (Goldstone & Medin, 1994), SME (Gentner, 1983), MAC/FAC (Gentner & Forbus, 1991), on precisely the problem of assuming away the object language by pre-defining or giving away to the program the representational scheme or relevant features in which to solve the problem. BACON, for example, quickly solves Kepler's problem with a precise, tabular representation of the solar system showing a primary body (Sun), a satellite body (planet), a time *T*, the two objects are observed, and two dependent variables—the distance *D* between primary and satellite, and the angle *A* found by using the fixed star and the planet as the end-points and the primary body (Sun) as the pivot point. Kepler, French notes, took 13 years (even more than Piaget's bead-children!) to sift through the data and flawed concepts of the solar system to find the relevant features. (How easy the tunnel–bead odd–even problem should have been!) Yet SME, he notes, uses an entirely different representation of the solar system, exactly suited to its programmatic purpose, to find an analogy to the Rutherford atom (see also footnote 8).

The ideal object language, or the 'representation module' as French terms it, which serves up the right representation from its store of all possible representations, will never be forthcoming. The features/representation of even a lowly credit card, if we wanted to pre-define it for *any* unspecified future problem, would be elusive, completely context dependent. French considers all the things a credit card (and thus the representation thereof) can be 'like:'

- A credit card is like a door key. (Problem: Motel door opening.)

- A credit card is like a breeze. (Problem: Need a fan.)
- A credit card is like a ruler. (Problem: Need to draw a straight line.)
- A credit card is like a catapult. (Problem: Need to launch a pea.)

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. 94).

And so we are back to experience. This is another form of the 'rating games.' When we rated 'knives as spoons,' it was described as the projection of the transformational dynamics of an invariance structure upon a possible component. So too when we rate credit cards as spoons (under a stirring transformation) or as fans (under breeze generation), etc. The list is endless because the card is fair game for an infinity of transformations under which new structural invariants can emerge. More precisely, it is fair game for insertion into an infinity of invariance structures. As earlier noted, there truly is no such thing as a purely abstract 'transformation.'

Let us consider one more variant of the approach, in this case Freeman and Newell's (1971) description of 'functional reasoning' in the design process. Designers, they argued, possess a knowledge of structures and their functional capacities. The process of designing a given structure proceeds by building up components through matching 'functional connections' of component substructures. The matching process is assumed to proceed in a heuristic fashion exactly analogous to the means–ends meta-framework of GPS. The object language, in this case, is the set of *functional requirements* and *functional provisions* characterizing each structure. To design a knife, using their example, we stipulate as part of the object language that a blade *requires* holding and *provides* cutting, while a handle requires being held and *provides* both holding of a narrow object and hold-ability by the human hand. It is possible to construct the knife due to the existence of the functional connection between blade and handle, and by the fact that the function we wish to provide, i.e., cutting, is not 'consumed' by the functional connection. The final 'goal' then for such a procedure is the 'provision of a cutting instrument hold-able by the human hand,' and the initial state is the list of structures with their functional provisions/requirements available for use. Again, the operation of the program can be considered the production of a proof that such an object can be constructed within the object language.

But let us think differently. Suppose we are asked to design a mousetrap. We are provided with several components: a piece of cheese, rubber bands, a 12 in. cubical box, pencils, a razor blade, toothpicks (the strong kind), a rubber eraser, string, tacks. What is the object language here? The simulationist must determine all the functional provisions and requirements of each—for he knows not (or should not know) what problem the program will be called upon to solve. But this is a lost cause. These are totally context-dependent. What is the exhaustive set of functional requirements/provisions of a pencil? These will emerge quite dynamically we shall see. At least, more realistically, he might attempt to define all the 'features' of each object. But what actually happens in thought? Our aim may be killing the creature, but there is no abstract transformation of 'killing.' Killing *is* an invariance across concrete forms of killing. So perhaps I contemplate crossbow shooting. This again places the potential components within a dynamic transformational structure. The stretchability and force of the rubber bands emerges, the sharpness and straightness of the pencils, the 'anchoring' potential of the side of the box to which I will tack the rubber bands, etc. These features become the object language I could provide the program. Or, contemplating beheading by axe, the length and requisite strength of the pencil emerges. I can groove the pencil and wedge the razorblade in to make an axe. The 'container' property of the box corner emerges, as I can prop the raised pencil-axe in the corner, a toothpick will prop it up, a rubber band tied to the pencil and tacked to the 'anchoring' feature of the floor will provide downward force, etc.

These 'features' of the objects dynamically emerge as a function of the transformations placed upon them via the invariance structure and the constraints naturally specified by the proposed structure, e.g., the crossbow requires anchoring points for the bowstring—the rigidity of the box can provide these. They cannot be all pre-set. New ones will always emerge. This is the problem, the simulation approach would hold that the features define or determine the analogy. In truth, *the analogy defines the features* (Robbins, 1976; Indurkyha, 1999; Dietrich, 2000). What is required is a device which can support the time-extended, dynamic transformations of experience.

Defining the object language, which is equivalent, as Narasimhan well knew, to creating semantic theories of the environment, is thus *the* problem. As Brian Smith (*On the Origin of Objects*, 1996) has brilliantly discussed, and Bergson long ago realized, there are no 'objects' in the universal field. To Smith, the definition or abstraction of these is the process of 'registration.' To Gibson, this is the process of the 'registration' or abstraction of invariance defined over the ambient energy flux and related to the action capabilities of the organism. The brain is presented a scale-less, undifferentiated field—from this it must define a scale of time and partition a world of objects, their motions and their interactions based upon the invariance laws it can abstract. This becomes its 'object language.' When this work is *pre*-accomplished, and inherently it must be so for the simulationist unless he has a device that can detect and 'store' invariance defined over time, we descend, as Narasimhan noted, unto the realm of mere proof procedures within a fixed theory—a fixed theory of the world defined by the object language. In other words, we deal in syntax.

4.2. The computational and beyond

It has been little remarked that the 'non-computational' thought of Penrose (1994), which he felt demanded conscious awareness, rests upon time-extended transformations defining invariance. Consider the proof that successive sums of hexagonal numbers are always a cubical number (hence a computation that does not stop). He has us imagine building up any cube by successively stacking three-faced arrangements that comprise hexagons—a back, a side, and a ceiling—giving each time an ever larger cube (Fig. 6). This is a dynamic transformation over time, in fact multiple transformations defining invariance. We can expand the hexagonal structures successively, from one, to seven, to 19, etc., each time preserving the visual hexagonal invariant. Then,

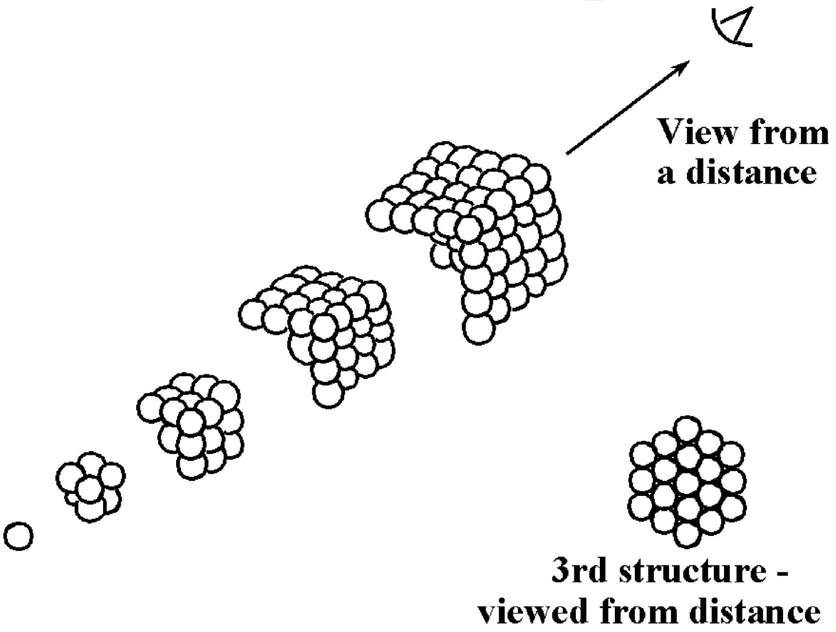

Fig. 6. Successive cubes built from side, wall, and ceiling. Each side, wall, and ceiling structure make a hexagonal number. (Adapted from Penrose, 1994.)

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.

As another example, he notes (Penrose, 1994) that if we consider an elementary fact of arithmetic, namely that given any two natural numbers a and b (i.e., non-negative whole numbers 0, 1, 2, 3, . . .), we have the property that

$$a \times b = b \times a.$$

Consider the case where $a = 3$, $b = 5$. Each side of the equation is different, and the two different groupings expressed can be displayed visually as

$$\begin{array}{ll} a \times b & (\cdot\cdot\cdot\cdot\cdot)(\cdot\cdot\cdot\cdot\cdot)(\cdot\cdot\cdot\cdot\cdot) \\ b \times a & (\cdot\cdot\cdot)(\cdot\cdot\cdot)(\cdot\cdot\cdot)(\cdot\cdot\cdot)(\cdot\cdot\cdot). \end{array}$$

A computational procedure to ascertain the equality of $a \times b$ and $b \times a$ would now involve counting the elements in each group to see that we have 15 in each. But we can see this equality must be true by visualizing the array:

$$\begin{array}{ccccc} \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot . \end{array}$$

If we rotate this through a right angle in our mind's eye, we can see that nothing has changed—the new 5×3 array we see has the same number of elements as the 3×5 array pictured. We see here, as in the case of the cubes, that the thing to which Penrose gravitates as a natural exemplar of non-computational thought is the perception of invariance. These perceived invariants form his 'obvious understandings' that become the building blocks for mathematical proofs. As we have seen of invariants, these obvious understandings, Penrose felt, are inexhaustible. From this he argued in effect, will arise the elements of an object language employed in a proof. But in this he was well preceded by the likes of Wertheimer (*Productive Thinking*, 1945), Arnheim (*Visual Thinking*, 1969), Bruner (*Beyond the Information Given*, 1973), Montessori (e.g., her mathematical program), Hanson (1958, 1970), and if one looks closely, Piaget (1946), and others.

Wertheimer (1945) described a visit to a classroom of children learning how to compute the area of a parallelogram. The teaching followed the traditional method of dropping perpendiculars and extending the baseline, and the teacher gave the students several problems to work involving different sizes of parallelograms. Wertheimer then got up before the class, drew a rotated figure on the board, and asked the class to work out the area. Only a small minority of the class was able to solve the problem, some of the rest responding that, ''they had not had that yet.'' Implicit in Wertheimer's discussion of the incident was the purely mechanical, 'human computer-like' knowledge the children had obtained. It went without saying that this was a degenerate form of knowledge in his opinion. It did not compare to the five-year-old he observed 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. Nor did it compare to the dynamic transformation exhibited by a five-year-old child who formed the cardboard parallelogram into a cylinder, then asked for a scissors to cut it in half, announcing it would now make a rectangle.

Yet, as Copeland (2000) has emphasized, Turing specifically defined the form of computation that he would formalize in terms of *mechanical* operations. He was thinking of the ubiquitous types of computation then found everywhere—the calculations of a bank officer balancing the ledger or of a clerk computing a total cost of purchase. 'Computation' consisted of the steps a human computer could carry out, a human acting mechanically *without intelligence*, i.e., *without semantics*. It was this form of computation that he would formalize in terms of the Turing machine.

As we have viewed the form and nature of the understanding underlying that which we can term a *semantic* 'computation,' it is clear that the Turing concept of computation is purely derivative. By this I mean that computation, in the Turing sense, is a simply a residue, in truth a spatialized husk of far more powerful operations of mind supporting representative thought, in turn based in the indivisible motion of the universal field. In a word, Turing computation is again a limiting case, fundamentally based in the 'projection frame' of the ever underlying abstract space and abstract time in which we tend to think (and theorize), itself a derivative concept from perception and its 'objects.' As with physics, this frame is what must be peeled away. It should not be a question then whether this narrow computation can account for the visualization transformations (experience) of Penrose. This entire paper, in the context of examining experience, has been devoted to the thesis that it cannot. The real question is whence the origin of this very notion of computation.

While Sejnowski and Churchland (1992) focus on the centrality of the function in the computational framework, yet as Copeland further notes, Turing, recognizing the limitation to his definition, reserved his 'O-machines' for functions beyond the computing power of the Turing machine. In truth, we could call the entire brain an O-machine—its input the holographic field, its output virtual action. But delineating (on a less macro order) the brain's computation in this far broader, dynamical sense, if indeed necessary, is beyond scope here.[13]

5. Conclusion

Let me pick up the thread begun in the Introduction, namely the obstacle posed by the Turing Test for any device incapable of experience. Hopefully the discussion of the form of device necessary to support experience has made it clear why current notions of computation face a difficult paddle against the wind. Semantics requires a device employing a far broader form of computation.

The syntactic device implies a world of discrete objects with fixed attributes. The motion of these objects is defined by discrete states. This world of abstract space and time fails completely to support qualitative invariance defined over continuous or melodic time. It fails to support invariants defined only over time-extended transformations. It fails to address the dynamic emergence of object attributes over the multiplicity of transformations in which they can be embedded. It fails to support the emergence of features under the transformation effected by analogy. In general, it fails to support the invariance structure of events. It fails to account for or incorporate the time-scaling of the perceived world which fundamentally supports the qualitative aspect of the world of experience. Finally, because it *does* imply the reality of abstract space, it fails to support the continuity of time-extended experience which is the essence of consciousness, and it fails to support the fundamental relationship of subject and object in terms of time essential to solving the problem of symbol grounding inherent within the syntactic approach.

Syntax, whatever the implementation, is not sufficient for semantics. To believe it is sufficient arises from a confusion of abstract space and abstract time with real, concrete time, with the quality derivative from this real motion, and the invariance existing throughout. This is the stuff of experience.

The framework presented here clearly rests upon an hypothesis, namely that the dynamics of the brain indeed support a modulated reconstructive wave within a holographic field. This in itself is a subject for future proof, and it is certainly not trivial. It has other challenges—a major one was highlighted in the discussion of voluntary action and the operative role of the image. Another was highlighted in the context of Piaget with respect to the work needed for understanding the process of dynamically embedded compositionality (or the Piagetian 'groups' such as INRC, etc.). Both these problems combine in the voluntary control dynamics supporting a thought process such as that underlying the Penrose cubes, the Piagetian semi-rotations (and underlying group), or visualizing the course changes to the next portage. But the Bergson/Gibson 'hypothesis' offers a viable theory of conscious perception, a basis without which any theory of cognition and memory has hitherto only been tentative. It truly acknowledges time. In fact, with respect to the problems of a scale of time, it raised questions 100 years ago that still have not dawned upon representationalism. It relies, integrally, upon dynamics, yet makes clear why dynamics must enlarge its vision of time before it

---

[13] Hoffman (1984), for example, insists that the neural process supporting his Lie transformation groups underlying perception is not computational. All of Yasue et al. (1991) can be taken as such, and many others—the work of Grossberg, Turvey, etc. There is a question whether, in trying to expand and define even 'broad' computation, we are attempting to stuff toothpaste back into the tube. Are we better off rejoining the natural sciences as Searle (1994) argued, and simply recognize that we are describing dynamics—like physicists describe the motion of planets or, yes, AC motors?

can support a theory of mind. And for added measure, it supports the capacity dear to the heart of the CH and representationalism—representative thought. It is, however, a different canoe to paddle.

6. Uncited references

Forbus et al., 1995; Searle, 1992

Acknowledgements

I would like to thank Robert E. Shaw, who at the inception of these thoughts, 25+ years ago, was the source of stimulation for many concepts, e.g., the possibility of the holographic field, the brain as a possible modulation device, invariance laws holding across time scales. All this supported the insights into Bergson and more. However, all aberrations here are indubitably mine. I would also like to thank one of the anonymous reviewers for many detailed questions and comments.

References

Adams, F., Aizawa, K., & Fuller, G. (1992). Rules in programming languages and networks. In Dinsmore, J. (Ed.), *The symbolic and connectionist paradigms*. NJ: Erlbaum.

Arnheim, R. (1969). *Visual thinking*. London: Faber and Faber.

Barsalou, L. W. (1993). Flexibility, structure and linguistic vagary in concepts: manifestations of a compositional system of perceptual symbols. In Collins, A., Gathercole, S., Conway, M., & Morris, P. (Eds.), *Theories of memory*. NJ: Erlbaum.

Bechtel, W. (1998). Representations and cognitive explanations: assessing the dynamicist's challenge to cognitive science. *Cognitive Science*, *22*(3), 295–318.

Bell, J. S. (1987). Beables for quantum field theory. In Hiley, B. J., & Peat, F. D. (Eds.), *Quantum implications*. London: Routledge and Kegan-Paul.

Bergson, H. (1889). *Time and free will: an essay on the immediate data of consciousness*. London: George Allen and Unwin.

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

Bergson, H. (1902/1920). Intellectual effort. In *Mind-energy, lectures and essays*. London: MacMillan, 1920 (W. Carr, Trans.; originally published in Revue Philosophique, Jan., 1902).

Bickhard, M. H., & Richie, D. M. (1983). *On the nature of representation*. New York: Praeger.

Bickhard, M. H. (2000). Dynamic representing and representational dynamics. In Dietrich, E., & Markman, A. B. (Eds.), *Cognitive dynamics: conceptual and representational change in humans and machines*. NJ: Erlbaum.

Birren, J. (1974). Translations in gerontology—from lab to life. *American Psychologist*, *10*, 808–821.

Bohm, D. (1980). *Wholeness and the implicate order*. London: Routledge and Kegan-Paul.

Bohm, D. (1987). Hidden variables and the implicate order. In Hiley, B. J., & Peat, F. D. (Eds.), *Quantum implications*. London: Routledge and Kegan-Paul.

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

Bruner, J. S. (1973). *Beyond the information given*. New York: W.W. Norton.

Cassirer, E. (1957). *The philosophy of symbolic forms*, The phenomenology of knowledge, vol. 3. New Haven: Yale University Press.

Chalmers, D. J., French, R. M., & 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.

Chalmers, D. J. (1996). *The conscious mind*. New York: Oxford University Press.

Chater, N. (1990). Connectionism and classical computation. *Behavioral and Brain Sciences*, *13*, 471–518.

Chopra, S. (1999). No cognition without representation? Dynamical computation and the emulation. *Cognet: Proceedings of NTCS99 module* (www.cognet.mit.edu/posters/NTCS99/Chopra.html).

Churchland, P. S., Ramachandran, V. S., & Sejnowski, T. J. (1994). A critique of pure vision. In Koch, C., & Davis, J. (Eds.), *Large-scale neuronal theories of the brain*. Cambridge: MIT Press.

Clark, A. (1997). *Being there: putting brain, body and world together again*. Cambridge: MIT Press.

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.

Cooper, L. F., & Tuthill, L. (1952). Time distortion in hypnosis and motor learning. *Journal of Psychology*, *34*, 67–76.

Cooper, L. F., & Erickson, M. (1954). *Time distortion in hypnosis: an experimental and clinical investigation*. Baltimore: Williams and Wilkins.

Craig, C. M., & Bootsma, R. J. (2000). Judging time to passage. In Grealy, M. A., & Thomson, J. A. (Eds.), *Studies in perception and action V*. NJ: Erlbaum.

Crowder, R. G. (1993). Systems and principles in memory theory: another critique of pure memory. In Collins, A., Gathercole, S., Conway, M., & Morris, P. (Eds.), *Theories of memory*. NJ: Erlbaum.

De Broglie, L. (1947/1969). The concepts of contemporary physics and Bergson's ideas on time and motion. In Gunter, P. A. Y. (Ed.), *Bergson and the evolution of physics*. University of Tennessee Press.

Dietrich, E. (2000). Analogy and conceptual change, or you can't step into the same mind twice. In Dietrich, E., & Markman, A.

B. (Eds.), *Cognitive dynamics: conceptual and representational change in humans and machines*. NJ: Erlbaum.

Dietrich, E., & Markman, A. B. (2000). Cognitive dynamics: computation and representation regained. In Dietrich, E., & Markman, A. B. (Eds.), *Cognitive dynamics: conceptual and representational change in humans and machines*. NJ: Erlbaum.

Edelman, G. M. (1989). *The remembered present: a biological theory of consciousness*. New York: Basic Books.

Falkenheimer, B., Forbus, K., & Gentner, D. (1989). The structure-mapping engine: algorithm and examples. *Artificial Intelligence*, *41*(1), 1–63.

Feynman, R. P., & Hibbs, A. R. (1965). *Quantum mechanics and path integrals*. New York: McGraw-Hill.

Fischer, R. (1966). Biological time. In Fraser, J. T. (Ed.), *The voices of time*. New York: Braziller.

Fodor, J., & Pylyshyn, Z. (1995). Connectionism and cognitive architecture. In Mac Donald, C., & Mac Donald, G. (Eds.), *Connectionism: debates on psychological explanation*. Oxford: Basil Blackwell.

Fodor, J. (1998). *Concepts: where cognitive science went wrong*. Oxford: Clarendon Press.

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 Riegler, A., Peshl, M., & von Stein, A. (Eds.), *Understanding representation in the cognitive sciences*. New York: Plenum.

Forbus, K., Gentner, D., & Law, K. (1995). MAC/FAC: a model of similarity based retrieval. *Cognitive Science*, *19*, 141–205.

Galton, F. (1883). *Inquiries into human faculty and its development*. London: Macmillan.

Garson, J. W. (1998). Chaotic emergence and the language of thought. *Philosophical Psychology*, *11*, 303–315.

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*(2), 155–170.

Gentner, D., & Forbus, K. (1991). MAC/FAC: a model of similarity based access and mapping. In Proceedings of the thirteenth annual conference of the cognitive science society. NJ: Erlbaum.

Gibson, J. J. (1950). *The perception of the visual world*. Boston: Houghton-Mifflin.

Gibson, J. J. (1966). *The senses considered as visual systems*. Boston: Houghton-Mifflin.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton-Mifflin.

Glenberg, A. M. (1997). What memory is for. *Behavioral and Brain Sciences*, *20*, 1–55.

Gleick, J. (1987). *Chaos: making a new science*. New York: Penguin.

Globus, G. G. (1995). *The post modern brain*. Philadelphia: John Benjamins.

Goldinger, S. (1998). Echoes of echoes? An episodic theory of lexical access. *Psychological Review*, *105*(2), 251–279.

Goldstone, R., & Medin, D. (1994). Interactive activation, similarity and mapping. In Holyhoak, K., & Barden, J. (Eds.), *Analogical connections*, *Advances in connectionist and neural computation theory*, vol. 2. NJ: Ablex.

Gray, J. A. (1995). The contents of consciousness: a neuropsychological conjecture. *Behavioral and Brain Sciences*, *18*, 659–722.

Grush, R. (1997a). An architecture of representation. *Philosophical Psychology*, *10*(1), 5–23.

Grush, R. (1997b). Yet another design for the brain? Review of Mind as Motion. *Philosophical Psychology*, *10*(2), 233–242.

Hanson, N. R. (1958). *Patterns of discovery*. Cambridge: University Press.

Hanson, N. R. (1970). A picture theory of theory meaning. In Radnur, M., & Winokur, S. (Eds.), *Analyses of theories and methods of physics and psychology*. Minneapolis: University of Minnesota Press.

Harnad, S. (1990). The symbol grounding problem. *Physica D*, *42*, 335–346.

Heidegger, M. (1927). *Being and time*. New York: Harper and Row.

Hintzman, D. L. (1986). Schema abstraction in a multiple-trace memory model. *Psychological Review*, *93*, 411–428.

Hoaglund, H. (1966). Some bio-chemical considerations of time. In Fraser, J. T. (Ed.), *The voices of time*. New York: Braziller.

Ho, Mae-Wan (1998). *The rainbow and the worm: the physics of organisms*. NJ: World Scientific.

Hoffman, W. C. (1984). Figural synthesis by vector fields: geometric neuropsychology. In Dodwell, P., & Caelli, T. (Eds.), *Figural synthesis*. NJ: Erlbaum.

Hoftstadter, D. (1995). *Fluid concepts and creative analogies*. New York: Basic.

Hurley, S. (1998). *Consciousness in action*. Cambridge: Harvard University Press.

Indurkyha, B. (1999). Creativity of metaphor in perceptual symbol systems. *Behavioral and Brain Sciences*, *22*, 621–622.

Ingerman, P. (1966). *A syntax-oriented translator*. New York: Academic Press.

James, W. (1890). Principles of psychology, vol. II. New York: Holt.

Jeannerod, M. (1994). The representing brain: neural correlates of motor intention and imagery. *Behavioral and Brain Sciences*, *17*, 187–245.

Jenkins, J. J., Wald, J., & Pittenger, J. B. (1978). *Apprehending pictorial events: an instance of psychological cohesion*, Minnesota studies of the philosophy of science, vol. 9.

Johnson, M. (1987). *The body in the mind: the bodily basis of reason and imagination*. Chicago: University of Chicago Press.

Keijzer, F. (1998). Doing without representations which specify what to do. *Philosophical Psychology*, *11*(3), 269–301.

Klein, D. B. (1970). *A history of scientific psychology*. New York: Basic Books.

Klinger, M. R., & Burton, P. R. (2000). Mechanisms of unconscious priming. I: Response competition, not spreading activation. *Journal of Experimental Psychology: Learniing, Memory and Cognition*, *26*(2), 441–456.

Kim, N., Turvey, M., & Carrelo, C. (1993). Optimal information about the severity of upcoming contacts. *Journal of Experimental Psychology: Human Perception and Performance*, *19*(1), 179–193.

Kock, W. E. (1969). *Lasers and holography*. New York: Doubleday–Anchor.

Kosslyn, S., & Koenig, O. (1992). *Wet mind: the new cognitive neural science*. New York: The Free Press.

Lakoff, G. (1987). *Women, fire, and dangerous things*. Chicago: University of Chicago Press.

Langley, P., Simon, H., Bradshaw, G., & Zytkow, J. (1987). *Scientific discovery: computational explorations of the creative process*. Cambridge, MA: MIT Press.

Mach, E. (1902). *Science of mechanics: a critical and historical account of its development*. Chicago: Open Court.

Marr, D. (1982). *Vision*. San Francisco: W.H. Freeman.

Marschark, M., Richman, C., Yuille, J., & Hunt, R. (1987). The role of imagery in memory: on shared and distinctive information. *Psychological Bulletin*, *102*, 28–41.

McCarthy, J., & Hayes, P. (1969). Some philosophical problems from the standpoint of artificial intelligence. In Meltzer, B., & Michie, D. (Eds.), *Machine intelligence*, vol. 4. Edinburgh, UK: Edinburgh University Press, pp. 463–502.

Minksy, M. (1975). A framework for representing knowledge. In Winston, P. (Ed.), *The psychology of computer vision*. New York: McGraw-Hill.

Mitchell, M., & Hofstadter, D. (1995). Perspectives on copycat. In Hofstadter, D. (Ed.), *Fluid concepts and creative analogies*. New York: Basic.

Morgenstern, L. (1996). The problem with solutions to the frame problem. In Ford, K. M., & Pylyshyn, Z. W. (Eds.), *The robot's dilemma revisited*. NJ: Ablex.

Nakamura, R. K., & Mishkin, M. (1980). Blindness in monkeys following non-visual cortical lesions. *Brain Research*, *188*, 572–577.

Nakamura, R. K., & Mishkin, M. (1982). Chronic blindness following non-visual lesions in monkeys: partial lesions and disconnection effects. *Society of Neuroscience Abstracts*, *8*, 812.

Narasimhan, R. (1969). Intelligence and artificial intelligence. *Computer Studies in the Humanities and Verbal Behavior, March*, 24–34.

Newell, A., & Simon, H. (Eds.), (1961). *Computer simulation of human thinking*. The Rand Corporation, p. 9.

Newell, A., & Simon, H. (1972). *Human problem solving*. Engelwood Cliffs, NJ: Prentice-Hall.

Norman, J. (in press). Two visual systems and two theories of perception: an attempt to reconcile the constructivist and ecological approaches. *Behavioral and Brain Sciences*.

Nottale, L. (1996). Scale relativity and fractal space–time: applications to quantum physics, cosmology and chaotic systems. *Chaos, Solitons and Fractals*, *7*, 877–938.

O'Regan, J. K., & Noë, A. (in press). A sensori-motor account of vision and visual consciousness. *Behavioral and Brain Sciences*.

Penrose, R. (1994). *Shadows of the mind*. Oxford: Oxford University Press.

Piaget, J. (1946). *The child's conception of movement and speed*. New York: Ballentine.

Pittenger, J. B., & Shaw, R. E. (1975). Aging faces as viscal elastic events: implications for a theory of non rigid shape perception. *Journal of Experimental Psychology: Human Perception and Performance*, *1*, 374–382.

Port, R., & van Gelder, T. (1995). *Mind as motion*. Cambridge: MIT Press.

Pribram, K. (1971). *Languages of the brain*. Engelwood Cliffs, NJ: Prentice-Hall.

Pribram, K. (1991). *Brain and perception*. NJ: Erlbaum.

Prinz, J., & Barsalou, L. (2000). Steering a course for embodied representation. In Dietrich, E., & Markman, A. B. (Eds.), *Cognitive dynamics: conceptual and representational change in humans and machines*. NJ: Erlbaum.

Pylyshyn, Z. (1973). What the mind's eye tells the mind's brain: a critique of mental imagery. *Psychological Bulletin*, *80*, 1–22.

Revonsuo, A. (1994). In search of the science of consciousness. In Revonsuo, A., & Kamppinen, M. (Eds.), *Consciousness in philosophy and cognitive neuroscience*. NJ: Erlbaum.

Robbins, S. E. (1976). Time and memory: the basis for a semantic-directed processor and its meaning for education. Doctoral dissertation, University of Minnesota.

Robbins, S. E. (2000). Bergson, perception and Gibson. *Journal of Consciousness Studies*, *7*(5), 23–45.

Robbins, S. E. (2001). Bergson's virtual action. In Riegler, A., Peschl, M., Edlinger, K., & Fleck, G. (Eds.), *Virtual reality: philosophical issues, cognitive foundations, technological implications*. Frankfurt: Peter Lang.

Robbins, S. E. (in press). Virtual action: O'Regan and Noë meet Bergson. *Behavioral and Brain Sciences*.

Rock, I. (1984). *Perception*. New York: Scientific American Books.

Schank, R. C. (1975). The structure of episodes in memory. In Bobrow, D. G., & Collins, A. (Eds.), *Representation and understanding*. New York: Academic Press.

Searle, J. R. (1980). Minds, brains and computer programs. *Behavioral and Brain Sciences*, *3*, 417–457.

Searle, J. R. (1992). *The rediscovery of the mind*. Cambridge, MA: MIT Press.

Searle, J. R. (1994). The problem of consciousness. In Revonsuo, A., & Kamppinen, M. (Eds.), *Consciousness in philosophy and cognitive neuroscience*. NJ: Erlbaum.

Searle, J. R. (1997). *The mystery of consciousness*. New York: New York Review of Books.

Semon, R. (1909/1923). *Mnemic psychology*. Concord, MA: George Allen and Unwin, B. Duffy, Trans.

Sejnowski, T., & Churchland, P. (1992). *The computational brain*. Cambridge: MIT Press.

Shaw, R. E., & McIntyre, M. (1974). The algoristic foundations of cognitive psychology. In Palermo, D., & Weimer, W. (Eds.), *Cognition and the symbolic processes*. NJ: Lawrence Erlbaum Associates.

Shaw, R. E., & Wilson, B. (1974). Generative conceptual knowledge. Unpublished manuscript, University of Minnesota.

Sherry, D., & Schacter, D. (1987). The evolution of multiple memory systems. *Pscyhological Review*, *94*(4), 439–454.

Smith, B. C. (1996). *On the origin of objects*. Cambridge: MIT Press.

Smolensky, P. (1995). Connectionism, constituency and the language of thought. In MacDonald, C., & MacDonald, G. (Eds.), *Connectionism*. Cambridge: Blackwell.

Sun, R. (2000). Symbol grounding: a new look at an old idea. *Philosophical Psychology*, *13*(2), 149–172.

Teng, N. Y. (2000). A cognitive analysis of the Chinese room argument. *Philosophical Psychology*, *13*(3), 314–324.

Turvey, M. (1977a). Preliminaries to a theory of action with references to vision. In Shaw, R. E., & Bransford, J. (Eds.), *Perceiving, acting and knowing*. NJ: Erlbaum.

Turvey, M. (1977b). Contrasting orientations to the theory of visual information processing. *Psychological Review*, *84*, 67–88.

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

van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, *21*, 615–665.

Vera, A., & Simon, H. (1993). Situated action: a symbolic interpretation. *Cognitive Science*, *17*, 7–48.

Verbrugge, R. (1977). Resemblances in language and perception. In Shaw, R. E., & Bransford, J. D. (Eds.), *Perceiving, acting and knowing*. NJ: Lawrence Erlbaum Associates.

Vicente, K. J., & Wang, J. H. (1998). An ecological theory of expertise effects in memory recall. *Psychological Review*, *105*(1), 33–57.

Vicente, K. J. (2000). Revisiting the constraint attunement hypothesis: reply to Ericsson, Patel and Kintsch, and Simon and Gobet. *Psychological Review*, *107*(3), 601–608.

Viviani, P., & Stucchi, N. (1992). Biological movements look uniform: evidence of motor-perceptual interactions. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 603–623.

Viviani, P., & Mounoud, P. (1990). Perceptuo-motor compatibility in pursuit tracking of two-dimensional movements. *Journal of Motor Behavior*, *22*, 407–443.

Wallace, J. G. (1956). Studies of perception in relation to age. *British Journal of Psychology*, *47*, 283–297.

Weiskrantz, L. (1997). *Consciousness lost and found*. New York.

Weiss, A. D. (1969). Sensory functions. In Birren, J. E. (Ed.), *Handbook of aging and the individual*. Chicago: University of Chicago Press.

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

Winograd, T., & Flores, F. (1987). *Understanding computers and cognition*. Norwood, NJ: Addison-Wesley.

Yasue, K., Jibu, M., & Pribram, K. H. (1991). A theory of non-local cortical processing in the brain. In Pribram, K. H. (Ed.), *Brain and perception*. NJ: Erlbaum.

Ziemke, T. (1999). Rethinking grounding. In Riegler, A., Peschl, M., & von Stein, A. (Eds.), *Understanding representation in the cognitive sciences*. New York: Plenum.