Analogical Reminding and the Storage of Experience:
The Paradox of Hofstadter-Sander

Phenomenology and the Cognitive Sciences, 2017, 16(3), 355-385

7/20/2015

Analogical Reminding and the Storage of Experience:
The Paradox of Hofstadter-Sander

Abstract

In their exhaustive study of the cognitive operation of analogy (*Surfaces and Essences*, 2013), Hofstadter and Sander arrive at a paradox: The creative and inexhaustible production of analogies in our thought must derive from a “reminding” operation based upon the availability of the detailed totality of our experience. Yet the authors see no way that our experience can be stored in the brain in such detail nor do they see how such detail could be accessed or retrieved such that the innumerable analogical remindings we experience can occur. Analogy creation, then, should not be possible. The intent here is to sharpen and deepen our understanding of the paradox, emphasizing its criticality. It will be shown that the retrieval problem has its origins in the failure of memory theory to recognize the actual dynamic structure of events (experience). This structure is comprised of invariance laws as per J. J. Gibson, and this event “invariance structure,” it is argued, is what supports Hofstadter and Sander’s missing mechanism for analogical reminding. Yet these structures of invariants, existing only over optical flows, auditory flows, haptic flows, etc., are equally difficult to imagine being stored in a static memory, and thus only exacerbate the problem of the storage of experience in the brain. A possible route to the solution of this dilemma, based in the radical model of Bergson, is also sketched.

Analogical Reminding and the Storage of Experience: The Paradox of Hofstadter-Sander

Introduction

"Analogy is the fuel and fire of thinking." So stated Hofstadter and Sander in their large tome on the subject of analogy (*Surfaces and Essences*, 2013). By this, the authors meant precisely that: "...without concepts there can be no thought, and without analogies there can be no concepts (pg. 3)." To elaborate further, they envisioned that, "…each concept in our mind owes its existence to a long succession of analogies made unconsciously over many years¸ initially giving birth to the concept and continuing to enrich it over the course of our lifetime (pg. 3)." In truth, Hofstadter and Sander see the scene of our thought as one of constant, voluminous, continual analogy. Analogy is the fundamental operation of thought. This fundamental operation is itself sub-served by an operation of memory, constantly dragging our past experience back into the present, ever placing the present event, via the context of the past, in an analogical light.

I will be accepting the thesis of Hofstadter and Sander (H&S) implicitly here. Their book lays out the supporting arguments in exquisite and extensive detail. Insights and discoveries abound within. They eviscerate current AI natural language translation programs, showing the source of the inadequacy in the inability to deal with analogy (pp. 367-383). But equally within is a profound problem, for as they unquestionably show, analogies are, of course, based upon our stored experience of the world. Yet as we shall see, Hofstadter and Sander find themselves struggling with this very concept, namely, the storage of experience in the brain. On the one hand, on the very basis of their extensive observation and insights, their analysis drives them to the position that only our retention of experience in all its vastness and detail can account for our power of making analogy. On the other hand, they cannot see how such vastness and detail can be stored/retained. In fact, they vote against it.

So here is a dilemma, or better, a paradox. If, to operate, analogy requires the retention of the mass of our experience in all its richness, and this mass of experience cannot be stored, how can this fundamental operation of thought work? Something must give. It is doubtful that what must give is the analysis of the nature of analogy and its requirement for the availability of the totality of our experience. Rather it may be that the origin of the dilemma resides in our notion, or perhaps better, our near-dogma, of the storage of experience in the brain.

This paper will be generally content to show two things: Firstly, an analysis of the dynamic structure of experience or of experienced events, exacerbates the "how can experience be stored?" problem. Secondly, the fundamental memory operation required for analogy – what has been termed *analogical reminding* – relies on this very dynamic structure (likely non-storable) in events. This is to say that this basic memory operation is itself intrinsically analogic, i.e., memory and analogy cannot be separated. Thus I intend to sharpen and clarify, and yes, bring to a point of criticality, the Hofstadter-Sander paradox. As the dilemma begs for an answer, however, I will at least sketch out a possible solution path based in Bergson's theory.

Analogical Reminding and Experience

An analogy that serves a salient purpose for Hofstadter and Sander features an event where the main author (DH) is visiting the North Rim of the Grand Canyon with his wife and young son, Danny. In the midst of this scene of great majesty and splendor, Danny bends down, intent on watching some insects in the sand. Many years later, when DH is visiting the great temple of Karnack, his friend, Dick, suddenly seems more focused on bending down and peering at a small bottle cap (an object for which he has an obsession for collection) on the ground. The former scene at the Grand Canyon comes rushing back to DH in memory. We have, in other words, an *analogical reminding*.

"Analogical reminding" was a term employed by AI theorist, Eric Dietrich (2000), as he contemplated the problem of analogy. Walking in an alley, he spots a jumble of garbage cans. Instantly he sees "Garbagehenge." In essence, the static structure of the garbage cans has driven recall of past experiences of Stonehenge, whether of visits or pictures, and in a meshing of these, the lowly garbage cans are now exalted in what we can call an "analogic" event. In this little example, Dietrich was struck by the fact that at the basis of analogy is an apparently instantaneous operation of memory retrieval that, of itself, is inherently analogic. In fact, he argued, it is this *analogic operation itself that is defining the features of the current event* (see also, Indurkyha, 1999). Currently the inverse is the case, i.e., it is held that the already-defined features of a current event drive or define an analogy to a past event with the same features (as for example, Gentner, 1983; Doumas, Hummel and Sandhofer, 2008; for a critique of the pre-defined features-aspect of these models, Chalmers, French and Hofstadter, 1992; French, 1999). Applying Dietrich's view in the Danny case, the retrieval of the Danny/Grand Canyon event and its superposition upon the Dick/Karnack event is itself defining the features of the analogy – attention on the trivial in the context of a grand geological/architectural feature. As an AI theorist, Dietrich acknowledged that he had no idea how such an operation could be implemented.

In the Danny analogy, the dilemma for Hofstadter and Sander thus begins. They now begin the search for the mechanism of analogical reminding, i.e., the basis for this fundamental operation of memory. The authors ponder how the original event was encoded such that the analogy with the future event is perceived (or the reminding caused). The notion of encoding (re Danny and the insects at the Grand Canyon) they note (p. 161 and ff.), is really how the event is being *conceptualized.* The term should not be taken as the "encoding" of the event in the brain as though storing a literal string – "Danny looking at trivial insects in the presence of a major geologic feature." It is actually the *category*, namely, "trivial sideshow more fascinating than the main event." Thus they stress, "…it is not the sequence of English words that we are talking about, but the abstract idea that it denotes (p. 171)."

The problem is twofold: Firstly, there is far more to the event than just this conceptual structure (if indeed this is encoded at the time, as opposed to Dietrich's dynamic creation via the analogy), and this "far more" is what the experience available to retrieval is. The conceptual structure is defined within the event, but what is the conceptual structure, in and of itself, without this event (and/or others) over which this structure can be defined? What is an abstract idea if not a pattern defined over or within concrete experience(s)? What is "a pure abstraction" (or "conceptualization")? "Bending," for example, involves something being bent, whether an elbow, a wire for a hanger, a branch for a bow, or Danny's body bending. Bending may be a concept, but as an abstraction, it does not exist without the concrete events of something being bent. Intrinsic to these bending events are very concrete forces and visual transformations.

Secondly, this "far more" to the event can support many other "encodings." A guide at Karnack, extolling its grandiosity, could have served as well to drive the retrieval of the Grand Canyon experience (perhaps also with a Grand Canyon-extolling guide) to include Danny's presence, or a sudden breeze stirring the desert sands at Karnack could retrieve the Grand Canyon experience which also could have had a sudden breeze, or perhaps a hawk soaring over Karnack serves for retrieval, or the conversational mention of buying hotdogs by the souvenir stand, or walking up a steep slope to the Karnack temple, or Danny's flapping shirt-tail, etc. Is there then a limit to the possible encodings?

Recognizing this latter underlying problem, they ponder the possibility of all experience, in detail, being stored in the brain:

> An experience would be captured in its entirety on our neurons, much as a film can be stored on a DVD. In the case of Danny and the Grand Canyon, having such a "total rote recording" would mean that the entire scene had been "filmed" in Doug's brain while he was experiencing it, and then that, some twenty years later, when he observed Dick stooping to pick up a bottle cap, this specific film had been reactivated in his brain by a mental search algorithm running through all filmed scenes in his memory (p. 172).

The impossibility of this vast storage of "event videos" is obvious, if only firstly because modern neuroscience can find nothing approaching an area in the brain that stores the replication of whole events, nor, for that matter, has anything like a molecular transcription process, acting in real time, been discovered. The authors further disparage the existence of anything like "pure visual resemblances" that could be used for matching events. Rather, they note:

> The moral is that we do not store in our memory a collection of "objective" events through which we run, seeking perceptual resemblances, whenever new events happen to us; rather the events that befall us get encoded – that is, *perceived, distilled and stored* – in terms of prior concepts that we have acquired p. 172, emphasis added).

But the authors are deeply perplexed on the subject of the very principles which govern the encoding of a concept for an event, i.e., just when, or yes, *if*, these abstract "essences" are indeed recorded as the "surface" of the event unrolls. Was the concept, "trivial sideshow…" really encoded as DH watched Danny long ago? They have already voted, as we have seen above, that this encoding has to exist, yet they are utterly uncomfortable with this very notion. For one thing, a reliance on such "encodings" of a previous event to explain a subsequent analogy seems very fortuitous. Secondly, as already noted, there is the question of how many concept encodings an event might have, for the event comprising the Grand Canyon visit is very doubtfully limited to supporting just one analogy. Thirdly, requiring previous encodings would appear to limit the emergence of analogies to precisely the set related to any previous encodings, yet the very "essence" of their book describes the inexhaustible, ubiquitous, continuous, creative, massive emergence of analogies via past experience.

To reinforce the last point, consider their discussion (pp. 189-190) of all the "categories" a glass can participate in. Categories, they argue (pp. 434-436) *are simply analogies*, discussed confusedly as though two separate things in the cognitive science literature, e.g., "Categories let people treat things as if they were familiar" (Spaulding and Murphy, 1996), versus, "Analogy is what allows us to see the novel as familiar" (Gick and Holyhoak, 1983). Just some in their "glass" categories are listed below:

- Artifact, industrial product, consumer article, fragile object, item of dishware, drinking glass, water glass, transparent object, recyclable object.
- Piece of freight, piece of merchandize, item for sale, unsold object, unsellable object, dust-gatherer.
- Glass for cold drinks, knickknack holder, toothbrush holder, home for tadpoles.

It is not necessary to engage in argument here as to whether this "categories as analogies" argument is definitive nor to enter into a discussion of concept formation – we are focused on the problem with finite encodings placed upon events versus the inexhaustiveness of analogical remindings. To note here then is that this exercise is reminiscent of French's (1990) considerations on the number of events that an object, say, a credit card, can participate in. In the context of the Turing Test, French proposed various tests for any computer attempting to masquerade as a human, where obtaining a passing grade relied totally on having the requisite concrete experience. One test was a rating game, with questions such as:

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

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

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

The list is endless. Says French (1999), "…no a priori property list for 'credit card,' short of all of our life experience could accommodate all possible utterances of the form, 'A credit card is like X' (p. 94)." Without the experience, he noted, one incurs the necessity of either pre-programming or training-up the association weights of all possible pairs of objects.

This "training-up of all possible pairs," it is worth noting, is exactly the thicket that current connectionist models of memory walk into when we consider our powers of analogy. The network of Rogers and McClelland (2004, 2008), for example, can be trained, given the input pair "SPOON CAN," to eventually respond with an output set such as STIR, SCOOP, LADLE, HOLD WATER. This finite set is the network's specifically trained "knowledge" or encodings on what a spoon "can do." Yet, when I imagine little brother launching a pea with his spoon at big sister across the table, I can easily see this addition to French's list:

- A spoon is like a catapult.

This (catapulting) is now something else that the spoon "can do," but the network, not having been trained on this response, would be incapable of recognizing it as a legitimate possibility (Robbins, 2008). Yet this derives naturally, as an analogy, from our experience of catapulting, whether by reading about Roman catapulting or watching "Pumpkin Chunkin" contests with pumpkins being flung by trebuchets.[1] For French, therefore, the number of

[1] One connectionist response is that there are better architectures that solve this problem. The Vector Symbolic Architecture (Gayler, 2003), offered as one example, is an effort to implement a more biologically plausible, more efficient architecture without using the troublesome backpropagation method of Doumas et al. (2008) or Rogers and McClelland (2008). Nevertheless, the fundamental view of analogy is exactly the same, i.e., analogy depends on a systematic substitution of components of compositional structures. In VSA, the systematicity and compositionality are considered the outcome of two operations, namely, binding and bundling, where binding associates "fillers" (Spoon, Coffee) with "roles" (Stirrer, Stirred), and where

"features" a credit card (or a spoon) can exhibit is inexhaustible. This is to say, equivalently, that we are unable to specify a limit to the number of concept "encodings."

French's examples will tend to be labeled "ad hoc" concepts (Barsalou, 1983), formed for the sake of specific goals, and thus making the emergence of analogy less continuous or ubiquitous than H&S argue. However, the existence of a distinction between ad hoc categories versus common (thus stored) categories, or, for that matter, between context invariant concepts versus context dependent concepts is being increasingly questioned (Barsalou, 1987; Barsalou, Wilson and Hasenkamp, 2010). Casasanto and Lupyan (2015), noting this trend, argue that *all* concepts are ad hoc, i.e., all are dynamically formed, via retrieval cues, in the moment. This means, as they argue, there are no concepts *ever* stored – all are being created in the moment via context-dependent retrieval of experience and no concept is ever quite the same from one time to the next. Thus, remembering the H&S position that concepts are in actuality dynamically formed analogies in the first place, each being, "a long succession of analogies made unconsciously…continually enriched over our lifetime," Casasanto and Lupyan reinforce this. But they diverge from the concepts-only-are-stored position to which H&S felt driven when considering whether experience in detail can be stored. Rather, they affirm an explicit reliance on continuous retrievals of past experience. The Danny-Karnack analogy would not be made via a previously encoded concept ("trivial sideshow…") but by the dynamic mesh of a past and present experience over which this concept emerges.

This evolving view, represented in Casasanto and Lupyan, reflects the position of this paper re H&S, namely, agreement on the continual emergence of analogies/concepts via memory retrieval of experience, but disagreement on reliance upon previously encoded concepts for this emergence. By implication as well, there is then the need for the detailed storage of all experience as H&S briefly entertained, but in their "event videos" considerations, rejected. Given the multiplicity of different analogies that can be derived from an event, it appears that we require the retention of the whole of any given event or experience. Put differently, if, as Dietrich envisioned, the retrieval of a past experience in conjunction with, or meshed with, a present experience dynamically creates an analogy within which features or concepts emerge, the experience as a whole must be available, for the same experience can serve in the creation of many other analogies via any of its multiple aspects.

The Abstractionist Barrier

But supporting the total recording of events in complete detail, at least per current cognitive science models, is very problematic. The origin of this dilemma begins with a huge lacuna in our understanding of the structure of events. This lacuna supports the notion, as Hofstadter and Sander stated, of events being "perceived, distilled and stored." This "distillation" concept expresses the theme of an encoding of an abstract structure of the event, or of "elements" or "features" of the whole event. Hofstadter and Sander are repeating the essence of a major position in the theory of memory, namely the abstractionist philosophy of memory storage (Cf. Crowder, 1993), i.e., perceptions inherently must be *reduced* and stored as schematizations – only selected elements or abstracted elements are stored. In Bar's (2007) version of the predictive coding framework (Hohwy, 2013, Clark, 2013), the reduction is termed the "gist" of events, where the brain can "encode in memory only a reduced, gist-based version of actual memories

---

bundling combines these role/filler bindings to produce larger structures. Therefore, the very elements of the composition, i.e., the features of objects, must be predefined (encoded), the values set, the relations fixed – all for the sake of allowing a syntactic process to unroll.

because these details can later be reconstructed with sufficient resolution" (p. 286). Barsalou (1993) is an excellent exemplar. Discussing the dynamic transformation of "biting" as in biting on a carrot, he thought this would be represented by (or stored as) three schematic states – "a mouth closed next to the object, followed by a mouth open, and then the mouth around the object" (p. 53).

There is no principled theory as to how or why the brain makes such a selection. *Which* states (or snapshots) of the biting face? Or in general, which elements of an event, and how are these identified in principle? In fact, the biting on a carrot is a flowing transformation – a facial flow field – which shares this flow-characteristic with other events we are about to examine. However, the abstractionist notion – the storage of static elements, a conceptual structure, a "distillation" – is ubiquitous and deeply engrained in theory.[2]

There is a structure prior to the conceptual, and upon which the conceptual could be (dynamically) built. It is not conceptual (as per H&S), nor mental, nor an encoding. It is a natural articulation of events. It is an intrinsic "encoding." In the next several sections that follow, we shall explore the basis for what I will term the *invariance structure* of events. Invariance laws exist over the optical flows, auditory flows and haptic flows (at minimum) that characterize events evolving over time. These laws are mathematical relations or structures intrinsic within these events. Again, they are not concepts in the sense considered by Hofstadter and Sanders, no more so than the law $a = f/m$ – intrinsic within an event where an object is accelerating while pushed by a hand – is a concept. These laws defining events, it will be argued, are the foundation of reminding. Ultimately then we will see the following:

a) We are dealing with invariants that exist only over these dynamic flows, i.e., they cannot exist in a static instant of time, or as a "bit" that can be translated over the nerves, or as a static bit stored in a neural memory.
b) The structure of invariants is immensely rich. As any component of the structure can sub-serve reminding, it is the entire time-extended, flowing event, in all its modalities – optical, acoustic, tactile – that must be (somehow) stored.
c) It is extremely difficult, however, to imagine these events – given that invariants existing only over dynamic flows define them – as being stored or "encoded" in a static memory as elements, features or sampled states.
d) It is this very structure of invariants – very problematically storable – that sub-serves redintegration, i.e., the basic memory operation of analogical reminding for which Hofstadter and Sander search.

The dynamically structured events we perceive comprise, of course, our experience. We cannot have a theory of how experience is stored unless we know what experience is. This is to say that we will require a theory of perception – the conscious perception of these events – as the

[2] There is clearly a form of "distillation." Paraphrasing Bergson (1896/1912), I practice Chopin's C# Minor Waltz a thousand times on the piano and the resultant is a motor memory that allows me to unroll the waltz at will. But this motoric resultant must be distinguished from the storage of, or memory of, each practice experience – the session when the teacher was mad at me, the session there was a terrible storm outside, and on. This is Tulving's (1972) episodic memory, or Sherry and Schacter's (1987) "System I" vs. "System II." This form of distillation (the motor or procedural memory) is a very different thing however from a memory based on the supposed breakdown and encoding of each of these practice experiences, by no principled method, into elements, features or time-slices for storage.

foundational prerequisite to any theory of memory. This is actually to say, we require a theory of the origin of our dynamically changing *image* of the external world. Without a model of perception and the origin of the image, then, as to how events are stored, we are merely throwing speculations into the wind. This fundamental fact is why – I am warning in preview – we will be led to consider Bergson's view on perception in the final sections of this paper.

The Invariance Structure of Events – Gradients and Flows

It was J.J. Gibson's (1950) insight that the surrounding environment contains a great deal of mathematical information useful to the brain. The information specifies the depth of objects, the form of objects, and how the body can act upon them. One such piece of information is the *optical flow field* (Figure 1). We see a flowing field like this routinely when we drive down the road in our car. The field is flowing by at the greatest velocity near the eye of the observer (or the driver of the car). There is a still point at the origin of the field called the "point of optical expansion" near the beginning of the mountains in the figure. This entire set or array of velocity vectors is called a "gradient," as a gradient changes gradually, in this case in terms of the values of the velocities, from zero to increasingly larger. This gradient happens to have a precise mathematical ratio – information useful to the brain. The velocity value of each vector is inversely proportional to the square of the distance from the eye, or $v \propto 1/d^2$.

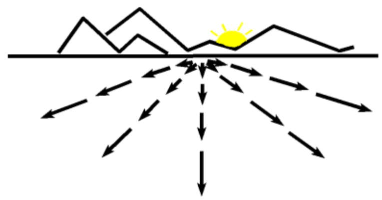

Figure 1. Optical flow field with its gradient of velocity vectors.

The actual surface of such a flowing field is also highly structured. We might be walking across a field of small rocks near the mountains, or a beach with its grains of sand, or a prairie with its grass, or across our kitchen floor with its tiles. These surfaces have what Gibson termed a *texture gradient* (Figure 2). Imagine the surface of Figure 2 as a rocky or gravel surface. The little circles (rocks in this case) are the "texture elements." The size of these elements and their horizontal separation (S) decreases in perfect mathematical proportion with the distance (d) from the eye or, $S \propto 1/d$. The vertical separation between each element decreases in proportion to the square of the distance, or $S \propto 1/d^2$. These texture gradients are ubiquitous – floors, beaches, lake surfaces, tiles, etc. Turn the gradients upside down – you see them as ceilings or the bottom of clouds.

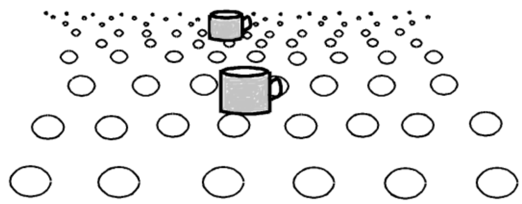

Figure 2. Texture density gradient (or tabletop) with a cup in two different positions.

Now we can imagine a coffee cup resting on the surface/gradient of a tiled table or patterned table cloth (Figure 2). In Figure 2, we can imagine that we have the *same* cup moving to two different positions – from back to front. The forward, larger view of the cup only occludes or covers two rows of (largish) texture units. The rear, smaller view of the cup occludes several layers or rows of smaller texture units. In fact, as the rear cup is moved to the forward position, there is preserved a constant ratio of the size of the cup to the number of texture units occluded. As the size (S) or height of the cup grows, the number of layers of units (N) it hides decreases, or $S \propto 1/N$. This is an invariance law. The ratio of cup height to texture units is an invariant – a ratio that does not change. It is this invariant that specifies the *size constancy* of the cup. The event of a cup moving towards us from one position to another is always experienced as a *cup of the same size* moving across the table – despite the growing size of the cup per se on the retina as the cup approaches. It is this invariant ratio, over time, that the brain is picking up that enables it to "specify" that there is a cup of constant size moving across the table.

Flows, Form and Time

More than the constant size of the cup, there is the form of the cup to account for as well. Form takes us further into the importance of invariants existing only over flows. Simultaneously an unreality emerges, namely that of "features" that exist in static instants of time, features that are thus thought storable in a static memory store. This notion is collapsing. An object recognition model, JIM (Hummel and Biederman, 1992) is illustrative of the static feature approach and its problems. The model described a neural or connectionist network for this recognition. In introducing the JIM model, Hummel and Biederman had argued that form recognition, for example, for a standard cube, cannot follow a straightforward model of decomposition into, and re-assembly from, sets of *features* – edges, vertices, straight lines. This is due to the fact that once the features are separately identified (disassembled), the actual spatial relations are lost and the features can now be recomposed in any number of possible ways where the results look not at all like the original cube (Figure 3). To defeat this problem, they moved to the notion of "geons." These are elementary solids such as cones, wedges, cylinders or bricks which are recognized by features such as straight or curved contours/cross-sections, and into which the forms are analyzed and which support the original spatial relations.

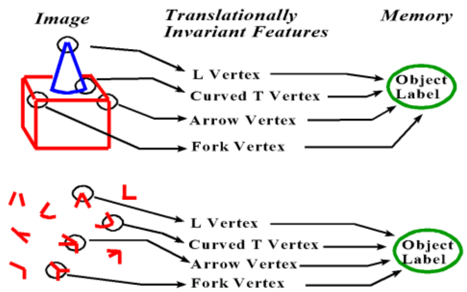

Figure 3. Both "objects" have the same features, and therefore would be "recognized" on the basis of a feature match. (After Hummel and Biederman, 1992)

In a review of the problem of dynamic form (Robbins, 2004), I noted the approach suffered on several points. For one, the classic fields to which neurons are thought sensitive, discovered by Hubel and Wiesel (1959; 1978), with their implication of detected, static features, simply cannot be regarded as the building blocks of a scene (Nakayama, 1998). Another, by their own admission, is that the "fast-enabling links" in neural structure and needed for the model have yet to be found. Worse, the model is subject to the *correspondence* problem. The correspondence problem is a problem inherent in assuming that the visual system matches corresponding points or features in successive frames of an event. It is deemed intractable. If Hummel and Biederman's example cube-cone were rotating, a model would not only have to deal with the feature jumbling, it would have to explain how the features (edges, vertices) are tracked (their identity preserved) from frame to frame (or sample to sample, or snapshot to snapshot). This holds too for the features of the geons. In essence, the primary problem is the static conception of form. Related is yet another problem, namely the once held concept that the brain assumes the object is rigid, i.e., that is there is a rigidity *constraint* employed to compute form (such as a rotating cube) under dynamic transformations (Ulmann, 1979a, 1979b, 1984, 1986). This concept met its demise as well. The current conception of the derivation of form is very different; it is based on velocity flows.

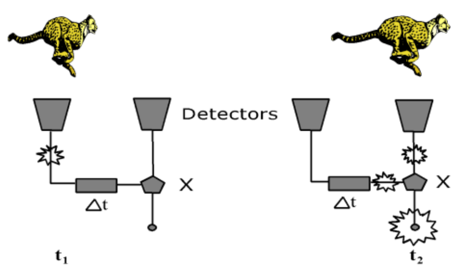

Figure 4. Reichardt filter or correlation model (Reichardt, 1959). It has two spatially separate detectors. The output of one of the detectors is delayed and then the two signals are multiplied. The output is tuned to speed. Many detectors tuned to different speeds are required for the true speed of a pattern, and the difference of pairs of detectors tuned to different directions is taken.

Driven by a desire to bypass the correspondence problem, current perception theory sees perceived form as derived from velocity flow fields in conjunction with Bayesian constraints. Adelson and Bergen (1985) described a general class of low-level models based on linear filters known as "energy models," initially developed by Watson and Ahumada (1983), for detecting the elements of dynamic form. These are addressed specifically to the detection of the direction and velocity of motion, for example, as an edge of our cube transits the visual field. They are an evolution from the correlation filter (Figure 4) of Reichardt (1959) for motion and speed detection, and there are significant formal connections to it.

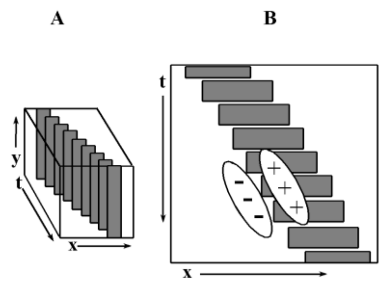

Figure 5. Motion as orientation in distance and time, (x,t). (A) is a spatiotemporal picture of a moving bar sampled in time. Velocity is proportional to the slant. (B) shows a spatiotemporally oriented receptive field that could detect the bar's motion. (Adapted from Adelson & Bergen, 1985).

In the correspondence problem, the position of a feature, say an edge, had to be tracked from frame to frame, and the distance change measured, to compute its velocity. The energy model does not extract position to compute motion. Motion is treated as spatiotemporal orientation (Figure 5), and the model consists of a network of "spatiotemporal filters" which respond to motion energy within particular spatiotemporal frequency bands. A network of these filters distributed across the visual field produces a net form of continuous output specifying the direction and velocity of motion of the edge.

This brings us to the model of Weiss, Simoncelli and Adelson (2002). A piece of background is yet in order. It should be understood that the receptive fields of the energy model filters are inherently "apertures," and thus the velocities of the flow cannot be estimated with certainty due to the limited view of each field. Figure 6 shows the problem. The card with the lines is moving to the right, so the card and its lines actually have a *horizontal* velocity. But when the card passes under the limited aperture, and the ends of the lines are obscured, only a downward motion of the lines is seen. The aperture forces an inherent uncertainty in measuring the actual velocity of the lines. More generally, this indicates that the visual system's measures of velocity are intrinsically *uncertain*. Therefore, the integration of a multitude of uncertain individual velocities must be inherently probabilistic. It is at this point of integration that Weiss et al. insert their fundamental, probabilistic (Bayesian) constraint, i.e., a probability estimate using a prior assumption about the nature of the world.

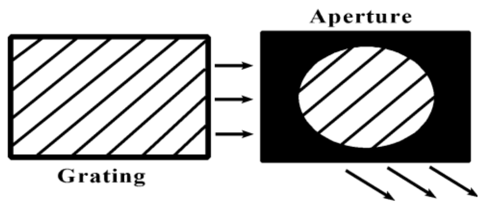

Figure 6. The aperture problem. The card with the grating is moving to the right, and passes beneath the card with the circular aperture. The ends of the moving lines are now obscured, and only the downward motion of the lines is seen in the aperture (Robbins, 2006).

The fundamental constraint used by Weiss, Simoncelli and Adelson (2002) in the resolution of these velocities is "motion is slow and smooth." This takes form simply as a mathematical constraint in their Bayesian model. (As I will be mentioning below at points the current "predictive" models of perception (e.g., Hohwy, 2013), it should noted that this *mathematical* constraint is far from a downstream-projected hypothesis that somehow looks like an expected event in the external world). The model explains a very large array of "illusions." In fact, due to this inherent measurement uncertainty, *all* perception, "veridical" or otherwise, the authors argue, must be viewed as an *optimal percept* based upon the best available information. Applied to the velocity fields defining a narrow rotating ellipse (Figure 7), for example, the violation of this "slow and smooth" constraint ends in specifying a non-rigid object if the motion is too fast (Mussati's illusion; Mussati, 1924). It is these constraints applied to the velocity flows, or their violation, that determine the rigidity of the form.

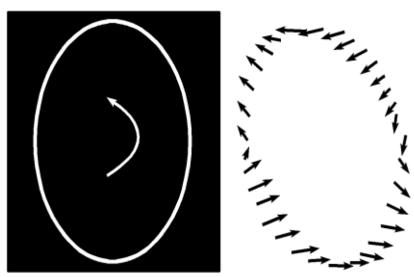

Figure 7. The normal velocity components (right) of the edge of a rotating ellipse (left). These tend to induce non-rigid motion. (After Weiss and Adelson, 1998)

If we were to consider a rotating "Gibsonian" cube, this form becomes a partitioned set of these velocity fields. As each side rotates into view, an expanding flow field (Figure 8) is defined. As the side rotates out of view, a contracting flow field is defined. The top of the cube is a radial flow field. The "edges" and "vertices" (i.e., "features") of this cube are now simply sharp discontinuities in, or junctures of, these flows. The implications of this are concretely displayed in a demonstration discussed by Shaw and McIntyre (1974) with a rotating wire-edged cube (Figure 9).

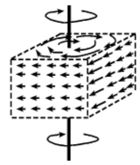

Figure 8. The Gibsonian Cube

A cube naturally has a symmetry period of four – it is carried into itself every 90-degree rotation. (Symmetry, it should be understood, *is* invariance). An equilateral triangle would be carried into itself or map precisely back upon itself every 120-degree rotation, or a symmetry period of three. When the rotating, wire-edged cube is strobed in phase with or at integral multiple of its symmetry period, it appears, indeed, as a cube in rotation. But when it is strobed out-of-phase, it becomes a distorted, wobbly, plastic or non-rigid object. In this wobbly "not-cube" case, the constraint (invariance) likely being violated via the arrhythmic strobe is this: a regular form displays a regular periodicity in time. The strobe is essentially taking snapshots of the cube. Yet these snapshots are not sufficient to specify the rigid cubical form we would expect; they are not sufficient to specify the straight lines, straight edges, corners or vertices – the standard static, geometric "features" of a cube.

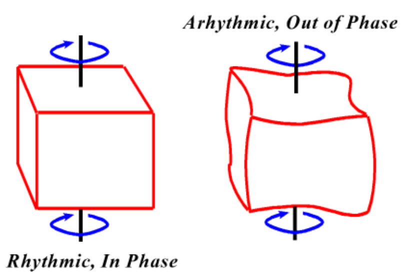

Figure 9. Rotating cubes, strobed in phase with, or out of phase with, the symmetry period (Robbins, 2004).

As Gibson long argued, the concepts of our Euclidean geometry – straight lines, curves, vertices, sets or families of forms related by geometrical transformations, even geons – while elegant, have little meaning to the brain, i.e., they are not the elements by which the brain constructs a world. Rather, the forms being specified are functions of the application of constraints on flowing fields (Gibson, 1966, 1979). The structure of the forms reflects invariants existing over these time-extended flows.

Flow Fields and Action

It is well known that the visual areas of the brain, V1 thru V5, are interconnected, modulating each other's processing ("re-entrant"). These visual areas themselves are modulated by connections from the motor areas. Indeed, a large number of recent findings have reinforced 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 (for example, Viviani and Stucchi, 1992; Viviani and Mounoud, 1990), while the appreciation of the importance to visual computation of re-entrant connections from motor areas to visual areas has also grown. We can begin here to consider the form of information in the light that "gets into" or is used by the action systems of the body.

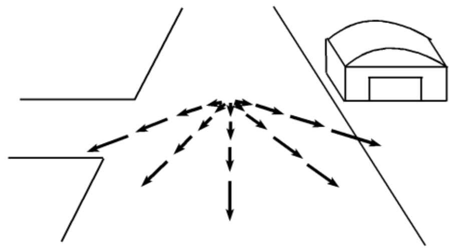

Figure 10. The Tau value defined over a flow field is used by a pilot to guide a plane.

There is a ratio, known as *tau*, that exists over flow fields (Kim, Turvey & Carello, 1993). It 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. For a bird coming in for a landing, this value specifies the *severity of impending impact*, and the bird can use it to modify his flight and create a soft landing. For a pilot, it is essential in landing a plane (Figure 10). Tau is an example of the "information in the light" that must be utilized by the action systems of the brain to specify possible (virtual) action and is integral part of the ongoing perception of an event.

Invariance in Dynamic Touch

Before we bring this all together in terms of the invariance structure of an event, let us spend a little time in another modality, that of touch. Kugler and Turvey (1987) consider the dynamics underlying a very simple exercise, namely where a person swings two rods, each with a weight at the end, in a pendulum-like motion using his wrists as the pivot points. The task is to swing the two pendulums comfortably together, at a common tempo, and 180° out of phase. In this case, the subject turns himself into a *virtual* single-wrist pendulum system. Such a system is characterized by a law derived by Huygens, wherein the center of oscillation (L_e) is defined as a function of the lengths (L) and masses (M) of the two individual pendulum systems:

$$L_e = (M_1L_1^2 + M_2L_2^2) / (M_1L_1 + M_2L_2).$$

The common period arrived at, while a natural period for the pair of systems *as a unit*, is not the natural period of either of the individual systems. The periodic times into which the person settles vary with the magnitudes of the pendulums, while the amplitude varies inversely with respect to mass and directly with respect to (rod) length. The information specifying the preferred steady states (dependent on the lengths and masses involved) is kinematic information based in the kinetics of muscle/tendon stress/strain distortions propagated globally via the nervous system as what ultimately can be termed *haptic flow fields*, i.e., flow fields defined over our systems for touch. This information serves as a source of constraint for the fields that produce it, and it guides the wrist-pendulum to its stable attractor state.

Kugler and Turvey, in their exhaustive analysis of the complex dynamics underlying this system, describe how change in periodic timing is brought about by change to internal parameters (stiffness, resting length, etc.) that preserve the amount of energy degraded as a constant over varying frequencies of oscillation. This is termed an *adiabatic* change (invariant), i.e., a transformation with the following ratio:

$$\frac{\text{Energy of Oscillation}}{\text{Frequency of Oscillation}} = k$$

I have spent a little space, time and depth here to impart the understanding of just how concrete, dynamic and *physical* (i.e., involving real forces) are the events and the invariance laws we are dealing with. We shall see that the simplest of events involve this form of information – information that is embedded in these very physical forces and fields. It is information, sorry to say, that is not detected (as per the Jeopardy-playing AI program, Watson) by scanning Wikipedia text.

We can consider another aspect of these "simple" forms of motion, for example simply "wielding" a metal rod a couple of feet long, or a short length of narrow board, or even "wielding" a spoon. Wielding is described under the concept of an "inertia tensor" (Turvey and Carello, 1995). A rigid object's moments of mass distribution constitute potentially relevant mechanical invariants since they specify the dynamics of the object. The object's mass (m) is the zeroth moment, while the first (static) moment is mass (m) times the distance (d) between the point of rotation and object's center of mass. The second moment is conceived as the object's resistance against angular acceleration. In three dimensions, this moment is a 3 x 3 matrix called the inertia tensor. The diagonal elements I_1, I_2, I_3, represent the object's resistance to angular acceleration with respect to a coordinate system of three principal axes (Kingma et al., 2004). There will be an inertia tensor (invariant), I_{ij}, specific to wielding the board and even wielding the spoon.

Stirring Coffee with a Spoon – an Invariance Structure

Now we consider a very simple, every day event, namely, stirring coffee with a spoon. The (ongoing) event has a time-extended *invariance structure*. We'll make this particular cup a cubical cup just to tie into the discussions above more clearly. Thus the swirling coffee surface is a radial flow field. The constant size of the cup, as one's head moves forward or backward or we move the cup towards or away from us, is specified, over time, by a constant ratio of height to the occluded texture units of the table surface gradient. As the coffee cup is moved over the table towards us, the tau value specifies time to contact and provides information for modulating the hand to grasp the cup (Savelsbergh et al., 1991). As the cup is cubical, its edges and vertices are sharp discontinuities in the velocity flows of its sides as the eyes saccade, where these flows specify, *over time*, the form of the cup. The periodic motion of the spoon is a haptic flow field that carries an adiabatic invariance – a constant ratio of energy of oscillation to frequency of oscillation. There will be an inertia tensor (invariant), I_{ij}, specific to spoon-stirring. All this and more is information specifying or structuring this simple event of spoon stirring.

Before we go too far, here is a definition of an invariance structure: An invariance structure can be defined as: *the transformations and invariants specifying an event and rendering it a virtual action*. The notion of "rendering it as a virtual action" has been given more depth elsewhere (Robbins, 2002, 2014), but already one can see, in the context of the τ value, how this applies. Let me re-emphasize here, these are *not* static structures. An invariant defined over a flowing field does not exist save only over the flow of the field; an invariant defined only over time does not exist, as Gibson noted, as a "bit" that can be transmitted along the nerves, i.e., it is not a static entity. All here is dynamic.

It is these structures, we can see, that are supporting the analogies discussed by French. When we rate “a knife as coffee stirrer,” we are *projecting the invariance structure of the event upon a possible component*. Can the knife support the forces required to create the motion in the liquid required in coffee stirring? When rating “a noodle as coffee stirrer,” the noodle comes off far less well. In rating a “sock as flyswatter,” we assess whether the sock, under this dynamic transformation, can attain the rigidity and mass necessary for ruining the fly’s day in our standard fly-swatting event. Rating the “spoon as catapult,” the spoon, under the dynamic catapulting transformation, now evidences new “features” or “properties” which support the invariance structure of the event.

In a seemingly tangential but very related subject, there is the frame problem, discussed heavily by AI for 30 years, never solved, then “faded” (cf. Wheeler, 2008). Thus we would consider how a hypothetical robot would detect that an event being observed contains features the robot would not expect to be occurring. We begin to see then that this is the form of knowledge – read, *experience* – that must be available to the robot. Assume she is performing coffee stirring and is checking her list of frame axioms – a set of axioms that in effect specify the features of the coffee stirring event and of the surrounding world that should remain the same: the size of the cup is unchanged, the spoon does not melt, the table does not collapse, the president remains the same, the sun stays in the sky… She is using these axioms to see if any happenings are unexpected. If the coffee surface is erupting in geysers, this is a violation of an invariance law specifying the event of stirring coffee, namely, the normal radial flow field. If the coffee feels like thick cement when stirring, this is a violation of the inertia tensor and adiabatic invariance normally involved. If the cup is bulging in and out, or growing up and down vertically, we have a violation of the normal size constancy of the cup. This form of knowledge is not going to be captured by the abstract symbolic rules of frame axioms.

As humans, we detect these violations instantly. They simply don’t resonate with our experience; rather, there is a dissonance with our experience. We are not searching, or we detect no time searching, a database (of frame axioms). This alone should be a warning that we may be an entirely different form of “device” than envisioned in the information processing framework with its frame axioms, themselves simply another form of encoding a finite list of event features. But in this subject, we have already touched on the fundamental operation by which events are retrieved and analogical reminding is supported. We turn to this next.

Redintegration and Invariance

We have, in essence, been searching for a model of redintegration for nearly 280 years, ever since Christian Wolff, a contemporary and disciple of Leibniz, first introduced the "law of redintegration" in his *Psychologia 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." A 1912 formulation of this by Dessoir (1912) stated: "Every idea tends to recall to the mind the total idea of which it is a part."

Examples of this phenomenon abound in everyday experience. Thus the sound of thunder may serve to redintegrate a childhood memory of the day one's house was struck by lightning. Perhaps, for example, we are walking down a road in the summertime and suddenly notice a slight rustling or motion in the grass along the embankment. Immediately, an experience returns in which a snake was encountered in a similar situation. 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.

Our difficulty has been in the fact that our models have approached these patterns as structures of static elements, and as though events can be treated as static snapshots composed of static features. The early mathematical models which were precursors the current connectionist framework, e.g., CHARM (Eich, 1985), TODAM (Murdock, 1982), are examples. CHARM proposed a model the paired associate task (e.g., pairing GRASS-SNAKE) using the concept that each item in the pair could be represented by a random vector of features. The vectors could simply be strings of 1's and 0's, where 1 stands for the presence of some feature, 0 for its absence. Thus GRASS might be represented as [1,0,1,0,1,1,0,0,1] and SNAKE (the paired word which must be learned) as [1,1,0,0,1,1,0,0,0]. The two vectors are multiplied together and the result stored in a "memory vector" M. If we call our two vectors A and B, then when A (GRASS) is presented again as a "cue," it retrieves a vector B'. If B' is close enough to the original B (SNAKE), in fact closer than any other possible response, e.g., E', D', F', we have a correct "remembrance." The event/patterns however are now simply vectors of 1/0 elements with, as Murdock (1982) noted, undefined psychological significance.

> Although one can make memory performance as good as desired by increasing N – the number of elements in the memory vector – what are these elements? If it takes 50, 100, 1000, or 10,000 elements to produce the necessary results, they are certainly not the cognitive features others have in mind (p. 625).

The connectionist models are no improvement. The earlier noted Rogers and McClelland (2008) model, for example, describes a network that given the input pair, ROBIN HAS (i.e., a pattern/vector of off units and on units), will elicit as responses, WINGS, FEATHERS, BEAK, and LEGS as opposed for TREE HAS, to BARK, BRANCHES, and LEAVES (other off/on unit patterns). They hold that this network is perfectly at home in the redintegrative situations of the ecological world. The input units can be construed as receiving perceptual input, for example the observation of a robin sitting on a branch, and the output units are predicting possible events or outcomes, say, the robin flying away. In a word, this is, as in all the above approaches, the simple connection or association of stored, static "elements" being used to explain memory. There is no attempt in these models to incorporate the invariance laws of dynamic events.

If, rather, we take to heart the actual structure of dynamic events, with the multiple invariance laws which describe them, we would be led to reformulate Wolff's law for redintegration as such:

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

In essence, redintegration implies that when a similar dynamic pattern, supporting a similar invariance structure, is evoked over the global state of the brain, the correspondent *experience* is reconstructed. I note that we are not describing a retrieval (and storage) mechanism here such as, for example, the connectionist paradigm seems to offer. What is being described here is the ultimate pattern or structure of events effective in the redintegrative operation, and which any mechanism must show that it can support.

In the traditional verbal learning context, this translates into very concrete implications that are obscured by the abstract focus of these experiments on words rather than on the actual events they express. One can propose an absurdly difficult paired-associate paradigm as far as verbal learning experiments are concerned (cf. Robbins, 2014). Call it the "A-B_i" paradigm. A list would look as follows:

SPOON-COFFEE
SPOON-BATTER
SPOON-OATMEAL
SPOON-BUTTER
SPOON-CORNFLAKES
SPOON-PEASOUP
SPOON-CATAPULT
SPOON-CHEESE
And so on…

It is absurdly hard since the stimulus words are exactly the same, the subject could have no clue what the appropriate response word is. But assume, in the spirit of the Subject Performed Tasks experiments (Zimmer et al., 2000), that the subject concretely acts out each event of the "absurd" list – stirring the coffee, stirring the batter, scooping/lifting the oatmeal, the cornflakes, cutting the cheese. The invariance structure of an event in effect implies a structure of constraints, and these constraints may be "parametrically" varied. Vicente and Wang (1998) alluded to this process in a different, more advanced memory context such as chess or baseball, as "constraint attunement." The (for example, blindfolded) subject may wield a Turvey type "tensor-object" in a circular motion within a liquid. The resistance of the liquid (a parameter value) may be appropriate to a thin liquid such as coffee or to a thicker medium such as the batter. The circular motion (a parameter value) may be appropriate to the spatial constraint defined by a cup or to the larger amplitude allowed by a bowl. The periodic motion may conform to the original adiabatic invariance (frequency/energy) within the event, or may diverge. We can predict that with sufficiently precise transformations and constraints on the motion of the spoon as a cue (either visual, or auditory or kinesthetic or combined), the entire list can be reconstructed, i.e., each event and associated response word.

The fact that our models must represent these event structures must be appreciated. It has, as just noted, predictive consequences, consequences which storing an event like coffee stirring in something like CHARM's static memory vector M, or in a connectionist network, cannot capture. An example is found in a demonstration by Jenkins, Wald and Pittenger (1978).[3] Capitalizing on the notion of the optical flow field, they showed subjects a series of slides that had been taken at fixed intervals as a cameraman walked across the university campus mall. Some slides, however, were purposely left out. Later, when subjects were shown various slides again and asked if they

[3] R*ecurrent* networks may be thought an answer to this. There are many variants of recurrent networks, but taking recurrence at its most elemental beginning, it is clear that if we begin with our standard network with its input layer, output layer and hidden (the middle) layers, and simply add a "context" layer of units which holds the previous state (or weight) values of the hidden units, this does not change the basic pattern – we still have only the learning of syntax rules (x is paired with b) and in a discrete state model of time. Recurrence, in and of itself, has nothing that intrinsically allows it to represent the dynamics of these events and the forms of invariance existing over the events' time-extension that we have discussed. If perchance recurrent nets can support wielding, stirring, and rotating/wobbly cubes, this form of encoding too needs to be explicated explicitly relative to how it supports the ecological laws of events. But in this, one cannot implicitly invoke a gratuitous source of continuity over these discrete states, e.g., a "consciousness," to support invariants that cannot exist in any one such state.

had seen the slide shown, they rejected easily any slide taken from a different perspective and which therefore did not share the same flow field invariant defined across the series. Slides not originally seen, but which fit the series were accepted as "having been seen" with high probability. But Jenkins et al. had created a "gap" in the original set shown to the subjects by leaving out a series of six continuous slides. Thus a portion of the transformation of the flow field was not specified. Subjects were quite easily able to identify these slides as "not seen." In this case, we are in effect varying parametric values defining a flow field, and more precisely here, varying the size of gaps in a flow to which we are sensitive. Other manipulations are possible, for example the slant of the gradient, the smoothness of the flow, the velocity of the flow, etc. In effect then, we again have a dynamic structure whose values can be parametrically varied, with definite consequences, one can propose, for things like familiarity judgments in recognition experiments. If a model, for example a connectionist net, cannot dynamically respond to these parametric variations, it is telling us that the models are not adequately representing events.

The original experiments of Pittenger and Shaw (1975), with their aging facial profiles generated via strain transformations on a cardioid (Figure 11), can be re-cast in this redintegration test framework. Originally the subjects looked at many pairs of faces, judging each time which of the pair was the older. Changing this to a memory task, a face of a certain age can be included in a set of various items successively presented to a subject. On the recognition task, a face transformed by a certain parametric aging value, i.e., a value of the strain transformation, is now presented. Familiarity values will be a function of the strain value of the transformation. The aging transformation works for animal faces too, even for Volkswagens – it can generate increasingly aging "Beetles." So we can have many different kinds of items in this test that eventually get aged (or un-aged) in a recognition phase.

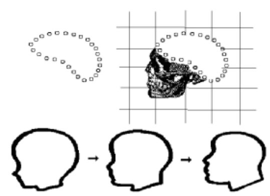

Figure 11. Aging of the facial profile. A cardioid is fitted to the skull and a strain transformation is applied. (Strain is equivalent to the stretching of the meshes of a coordinate system in all directions.) Shown are a few of the possible profiles generated. (Adapted from Pittenger & Shaw, 1975)

These are memory experiments that have been done or are waiting to be done. Unfortunately, they are but a small beachhead into the research program that should be carried out in this invariance structure framework. Such a program is yet neglected. The experiments provide a level of experimental control over memory stimuli and therefore the ability to tune performance that is the essence of scientific control. The experiments demonstrate that memory and its supporting brain dynamics are extremely sensitive to the invariance structure of events and the actual "parameter" values of the transformations involved. The importance of this is that any theory that claims to be a model of memory must be able to support this dynamic structure in events. Any model then that is "encoding" or reducing or "gisting" these events for storage in the brain should be required to answer these questions:

1) What is the principled method by which these events are reduced or "gisted" (and/or stored as static slices)?
2) How does this encoding, reduction or gisting capture the dynamic invariance structure that actually characterizes the events with their transformations?
3) In its retrieval performance, will the model support (show sensitivity to) parametric variations of these dynamic structures?

This point is apparently elusive. The predictive models, currently of great interest, are not particularly specific about, or worried about, exactly how a current event achieves its associations to (retrievals of) other past events. Bar (2011) for example, requires only that "associations" be activated. But in the prediction context of this predictive framework, where the brain is taken to be projecting a perception/event hypothesis neurally downstream upon an ongoing event, the questions above quickly apply, now to the hypotheses or projections themselves. Bar (2007), in this predictive context, notes Freyd's (1987) "representational momentum" experiments where for example a short series of frames of a rotating rectangle is presented, each frame showing the rectangle rotated a bit more. The subjects' memory test results show a memory for a rectangle in the next (not shown) position, for Bar, suggesting that the brain is generating an image of the next frame to be presented and supportive of the predictive notion. But these experiments are simply special cases of experiments such as that of Jenkins et al., noted above, with its flowing field down the campus, or for that matter, frames of coffee stirring. This implies, again, that the brain is responding to the dynamic structure of events, with invariants defined over time, and that again the predictive models must explain how these structures are faithfully encoded (for future use), to include the principled method of the reduction of these events to elements or gist. As these predictive models are actually *comparator* models, matching, in an ongoing event, a predictive or projected hypothesis on the form of the event against the actual ongoing event, this would include a far more concrete model of how such a comparison could actually work in real time for such a dynamically changing event. Imagine, for example, attempting to describe the origin of, and nature of, the (ever changing) projected hypothesis/image and as well the real time matching process (and at what scale – each $1/10^{th}$ second?) for the wobbly, plastically changing, not-cube with its flowing fields, or the coffee-stirring with its adiabatic invariants, inertial tensors, and flow fields.

Return to Analogical Reminding

Underlying analogical reminding, we can argue, is the operation of redintegration. Dietrich's lowly Stonehenge-like garbage cans were simply a "static" event. We have been dealing with events that are far more dynamic – continuous transformations or change in the external field. While stirring coffee in the morning, I am reminded of mixing cement for the barn wall yesterday. I come to breakfast in the morning, look outside, see the bird feeder swaying, slanting in the breeze, and immediately last night's dream returns in which there was a large ship run aground near the shore and listing heavily to its side. Or while stirring coffee, I am reminded of a co-worker stirring up trouble among my fellows at work.

This latter reminding – that of the trouble-stirring co-worker – surely strikes one as a mapping to a high abstraction. But to make a quick plug for the generality of what is being discussed here, these abstractions are nevertheless built upon very concrete events. Johnson (1987) and Lakoff (1987) have shown this foundation of metaphor and analogy (both being mappings) in basic events very effectively and in detail; they have simply not discussed these events in terms of the invariance structures necessarily defining them. They frame things, rather, only in terms of the bodily basis of the mapping. "Anger," Lakoff argues, involves a model of increased body heat, increased internal pressure (blood pressure, muscular pressure), agitation, and interference with accurate perception. It leads to metaphoric statements such as: "He makes my blood boil," "Simmer down," "Let him stew," or "The entire office was simmering." Boiling, simmering, or stewing themselves, however, are events with definite visual, auditory and tactile invariance structures. The basis for these mappings necessarily involves these dynamic event structures. When speaking of "Joe stirring up the office," we are relying on invariant elements in the stirring structure – a medium being stirred (emotions), an instrument doing the stirring (Joe's words/actions), a force applied (Joe's anger), elements in motion (people's emotions), etc. – but these invariant elements themselves only exist and have their meaning within the concrete structure of the event with its forces, frequencies, tensors, swirls, flow fields and more. "Joe is driving the point home," may be considered a common (dead) linguistic metaphor, requiring no special work at mapping, but this redintegrative mapping surely, perhaps more automatically, occurs, e.g., to the dynamic event of a hammer driving a nail with its forces, tensors, haptic and visual flows, etc. If not, the sentence makes no sense; the mapping is the basis of its meaning.[4]

Any pattern or sub-pattern, then, within the overall structure/pattern of a current event is capable of the redintegration of other events with a similar pattern. This includes patterns like Danny bending down to pick up a small object in the presence of a vast geological or architectural feature.

But these patterns are invariance structures. The swirling coffee surface is a flow field. Even the form of the cup is an invariance over flows, the periodic motion of the spoon is a haptic flow, the bending down is a flow. This brings us to the critical question: are such events, where "features" are only invariants over dynamic flows, actually stored in the brain?

The Problem of Storing Dynamic Events

How can a flow be stored in a static, three-dimensional memory, i.e., in the brain? A very natural solution is the concept that the brain is storing a series of samples of the ongoing event, as though we had a set of snapshots laid out upon a desk top. But the problems with this approach are legion. We have seen, for example in the case of the strobed rotating cube, that the brain is uninterested in samples. Each strobe can be considered to be providing a sample, yet the form of the cube is not specified in a single sample. For the brain to be employing samples in this case, it would need some form of precognition to adjust itself to the strobe rate such that it is always sampling in phase with (or at an integral multiple of) this particular rate. And what if there were two or three cubes rotating simultaneously at different rates?

[4] The "Joe stirred" example involves syntactic elements (Joe/subject, stir/verb, office/object…). Such elements on the surface seem "static" concepts, not reliant on flows. One simple thing to note, however, is that these very syntactic structures are also used in the "semantically deviant" sentence, a sentence which on the surface seems to have no meaning, e.g., "Joe stirred the building," or "The building stirred the coffee." Yet it is by adjusting and preserving the underlying dynamic invariance structure of the event that we provide even these sentences meaning. For example, we imagine that Joe is using a mini-Lego building to stir the coffee, where the tiny size of building-stirrer preserves the relative spatial relationship required between stirrer and the container/cup in a stirring event.

Further, storing samples introduces a regress. The samples are static, like the above-noted photos laid on the desktop. How then is the actual motion of the event registered? The brain would require some form of internal "scanner" sweeping across the photos. But now we have to answer the new question of how the scanner itself registers motion.

The physicist, Walter Elsasser (1987), noted in his thoughts on memory, "No transcription mechanism has ever been discovered that writes these flowing events in molecular structures in real time, let alone reads them back out." To my knowledge, this situation has not changed. Within this "transcription" notion, however, there is another negative implication. It is a standard conception, as we have seen, that the brain is parsing or reducing events (the samples) into features and storing them in separate locations in the brain. The rotating cube, for example, is parsed into its features – edges, vertices – and these features (within a snapshot) are stored. The features of the coffee stirring event, snapshot by snapshot, would be stored in separate spots.

As noted, there is no principled theory of the origin and definition (particularly in finite numbers) of these features of an event or, to say the same thing, of the "elements" of the event that are being stored. This lack should now not surprise us. What, one can ask, are the features of the coffee stirring event that can be stored as static elements? We are being asked to store an adiabatic invariant defined over dynamically changing physical forces over ongoing time. We would be storing the flow field of the swirling surface and the ongoing dynamics of the inertial tensor describing the wielding of the spoon. For the rotating cube, there are the discontinuities defined at the junctures of flow fields defined over the turning sides, the junctures themselves then existing only as invariants over dynamic change. But even if these elements or features of an event could be defined via some principled theory, there remains the fact that to account for the perception of the ongoing coffee stirring event, the information from the event (yet to be "perceived") must be parsed or disassembled into its features, each stored, then immediately retrieved from the various spots in the memory store, reassembled as the event, and somehow projected, whether internally or externally, as a frame of the past event – depending on your theory of perception – all in real time. Simultaneously the system is yet parsing another event frame or sample (yet to be perceived) of the ongoing event. This real time disassemble-reassemble process, never defined or explicated in any theory, and using "elements" with no principled origin, is what is implied in the conception that a short term or "immediate" or "iconic" memory supporting an ongoing perceived event (such as our "coffee stirring") simultaneously involves a parsing and storing of elements/features of the event. Note too that this process in effect relies on something like "event geons," for just as the form geons of Hummel and Biedermann served, in effect, as the handy picture on the top of the puzzle box that helps us put the pieces together properly, so too we would need something (a "coffee-stirring geon") to guide the reassembling of the event. A simple "index" linking the disparate stored elements is not sufficient for event reconstruction.

Unmentioned in the above discussion is an overarching problem, namely that the disassemble-reassemble of snapshots process appeals surreptitiously to a consciousness (or something – some form of "event glue") that is providing continuity such that an event flow is perceived. An event frame is disassembled into static features, reassembled, then "projected" as a perception, then another, then another… Each is only a static projection, a frozen moment in the event. There is no continuity of flow, no more than there is actual continuity in the successive projections of movie frames on a screen. Something else is providing the continuity, some other form of memory. In ignoring the source of this underlying continuity, we are begging the explanation of the perception of change.

But if storing samples/features of dynamic events is riddled with difficulties – and we have barely begun to go as deeply into the problems as we could – where does this leave us in terms of a theory of storage? In fact, where does this leave us on the perception of ongoing events – the "rotating" cube, the "twisting, falling" leaf – inherently based, as is perception itself, in a form of memory? There is little for us to find as an alternative in standard cognitive science. We are led back, full circle to Hofstadter's and Sander's dilemma, to wit, accounting for the storage of experience at all – experience with its spoons stirring coffee, faces biting carrots, leaves falling – let alone then accounting for our powers of analogical reminding.

The Picture Thus Far

What we have seen here is the following:

a) Upon examination, the structure of dynamic events, that is, our experience, involves a structure of invariance laws, and particularly, invariants defined over time, over dynamic forces and/or over flows, e.g., haptic flows, optical flows, etc.
b) The fundamental memory retrieval operation, namely, redintegration, relies on cue events which reinstantiate the invariance structure (or a subset thereof) defined over these flows. This operation is itself the basic mechanism underlying "analogical reminding" for which Hofstadter and Sander searched.
c) Nowhere in (a) or (b) is there any hint of "distillations," static "elements" of events or static features of these events being stored and used in the retrieval operation. There is only the description of invariance over dynamic flows – both over the whole of the original event and over the retrieval of that event.
d) The storage of flows and invariants defined only over flows requires a form of memory storage and retrieval obviously different from any contemplated in current cognitive science, based as it is upon the concept of static elements.

Hofstadter-Sander and their paradox bring us to contemplate analogical reminding and its dependence on the storage of our experience. Given what we have seen above, we can appreciate more fully the profound difficulties for our theory of memory that inevitably we find ourselves forced to face.

All of this, of course, begs a question: what might the answer to the storage of experience be? We will look at a sketch of a possible answer. The difficulties noted above force us, I believe, to consider the radical theory of memory found in Bergson (*Matter and Memory*, 1896), and this begins in our concept of just what experience is.

Experience – The Image of External World

We cannot have a theory of how experience is stored, I noted earlier, without knowing what experience is. We have seen in the foregoing that it consists of dynamically structured events with their invariance laws. But it is more basic than this. The coffee cup with spoon stirring away on the table before us is our ever transforming image of the external world. It is accounting for this image that has been the problem. If the brain is projecting a predictive hypothesis down neural pathways on the coffee cup with stirring spoon (Hohwy, 2013, Clark, 2013), what can the "hypothesis" actually look like? In the brain, we see nothing like an image of the world, we see neither coffee cup nor spoon; we see only, say, to pick a level of description, neural-chemical flows. In a computer, we again see only electrical impulses, bits flipping on and off. From the vantage point of an external observer, we can attribute to the brain or to the computer the existence of an image over these processes, but it is a mere attribution.

For the last twenty years, this problem has been incorporated within Chalmers' hard problem: how, given any neural or computer architecture, does one account for the qualia of the perceived world? But this formulation, in terms only of qualia, has been very misleading. For one thing, it has not been understood that form is qualia – the wobbly, plastic-like, not-quite-a-cube is a quality far different from the rotating, rigid cube (Robbins, 2013). Dynamically changing form as being itself, qualia – not just the "redness" or "blueness" of the rotating cube or of the wobbly cube – is exactly the intuition of Hardcastle as she enumerated her examples of qualia: "… the conductor waving her hands, the musicians concentrating, patrons shifting in their seats, and the curtains gently and ever-so-slightly waving" (1995, p. 1). The entire kitchen is qualia, to include the coffee cup, our stirring the liquid, all the forms therein – tables, chairs, the cup, the gently waving curtains – as well as the colors – the whiteness of the cup, the brown of the coffee. There is nothing in the image of the external world that is not qualia. Again, it is accounting for our image of the external world that is the more general problem.

This was Bergson's starting point in 1896. He noted that of course we see nothing like a photograph of the external world within the brain. But he went on:

> 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, emphasis added)

Here, in the photograph "already developed in the very heart of things and at all points in space," Bergson was stating, nearly 50 years Gabor's discovery and nearly 85 before Bohm's (1980) generalization, a prescient vision of the universe as a holographic field. In an earlier passage, he had previewed this, viewing the field as a vast interference pattern – a vast field of "real actions." Any given "object" acts upon all other objects in the field, and is in turn acted upon by all other objects. It is in fact obliged:

> ...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)

But Bergson envisioned the role of the brain quite differently from anything in current theory – it is not, for example, that the brain is simply a "hologram" as per Pribram (1971). Following hard on the "photograph" passage he noted:

> 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, emphasis added)

In essence, Bergson is viewing the brain, taken over all its physical dynamics, as creating a very concrete wave – a *modulated reconstructive wave* passing through the holographic field. As a reconstructive wave, it is specific to, or specifies a portion of the information in the field, and now by this process, an "image" of the field – the coffee cup on the table. As a reconstructive wave when modulated to a certain frequency and passing through a hologram selects a set of information from the hologram plate – one wave front from a set of stored wave fronts – specifying the source of that original wave front, so in the brain's case, the selection principle is via information relating to the body's ability to act. This is why Bergson stated, in essence, that perception is virtual action – we are seeing how we can act. Simultaneously, the information in the field that is relatable to action is, as we have seen, the dynamic invariance structure of the events in the field. We can say then that the modulation of the brain, as a reconstructive wave, is driven by the dynamic invariance structures of external events.

Critical in this picture is Bergson's model of time. In watching the spoon stirring, the coffee swirling or a fly buzzing by, we are viewing the past, i.e., a portion of a past transformation of the universal field. But Bergson viewed this transformation, i.e., "time," as indivisible. There are no mutually external "instants" in this transformation, where each "instant" passes away into non-existence (the "past") as the next instant – the "present" – arrives. As is well known, he viewed this transformation as a melody where each note (read, instant) interpenetrates the next forming an organic continuity, and where each note reflects the entire preceding series. But with this view of the transformation of the universal field, we now have the source or basis for a *primary memory* (to appropriate James's (1890) term, used now in an even more primitive mode). This primary memory property of the ever-transforming material field allows us to see stirring spoons, leaves twisting and falling, flies buzzing by, etc., for the brain's specification then is simply to a portion of this (past) indivisible motion of the field. It is a motion that does not and has not fallen into the past (therefore ceasing to exist) and thus has to be stored somehow in the brain (we have seen some of the difficulties above) to preserve it.

What this says is that perception – our experience – is not occurring solely within the brain. Rather, it is a specification of events, an optimal specification that is at a scale of time determined by the brain's dynamics – a "buzzing" fly versus a fly slowly flapping his wings like a heron, a plastically changing non-cube vs a rigid, rotating cube. These events are precisely where they appear – external, within the ever transforming external field. Therefore, since experience is not occurring solely within the brain, experience cannot be solely stored there in the first place. This, I would propose, is why current memory theory is having such trouble determining how the brain in fact stores experience. The question thus becomes how this experience is retrieved, a process we have already seen will rely on the afore-discussed principles of redintegration.[5]

Not Storing Experience

The indivisible flow of experience creates a time-extended structure of being. Bergson saw the brain, embedded as the leading edge of this ever growing structure, as a sort of "valve" which allows experiences from the past to enter (or equally, be shut out). The valve is opened in various

[5] With this view of time as indivisible flow, the brain is also a four-dimensional, time-extended object. Hence, at least the neural aspects of any external event being specified can be considered "stored" in this sense, i.e., the neural flow of processes shares this time-extension. The practical consequences of this, however, I have given little consideration.

forms coordinate with the current action state of the body, or as he would term it, the "nascent" or possible actions being prepared or sketched. In a word, the momentary configuration of Bergson's "valve" is very much a function of the invariance structure of the external event and thus the form of modulation of the brain being driven by this structure. If I am a pilot coming in for a landing, the runway flow field, via the tau ratio, is carrying information for the adjustment of my action, both ongoing and virtual. And in general, we return to the fundamental law of redintegration described earlier, remembering the definition of invariance structures with respect to virtual action, wherein a present event redintegrates a past event or set of events sharing a similar invariance structure. But it should be noted simultaneously that this is why the view that amnesias or aphasias prove that experience is stored in the brain is arguable at best. The fact that a konk on the head appears to destroy stored experience can be equally taken to be indicating that complex mechanisms and structure allowing the brain to support this modulation or redintegrative process, and/or explicit recall process, have been damaged.

It is easy to speak of this redintegrative process in terms of the brain, now again acting as a reconstructive wave, reconstructing past events as though we are again dealing with a wave traversing a hologram, i.e., retrieving events from the holographic field taken as a four-dimensional structure. This is not what is happening however. In perception, the coffee cup with the stirring spoon or the fly buzzing by are indeed optimal specifications – now images – of the past transformation of the external field, a field that as a vast interference pattern is in fact non-image-able and wherein images are only yielded via this reconstructive wave process. As dynamically changing images (experience), sharing equally in the indivisible transformation of the field where there are no instants falling away into non-existence, this changing image/experience too is a non-erasable, persistent reality within the whole ongoing transformation of the field. In Bergson's terms, these images of the past are now "virtual." When we encounter in the present yet another coffee stirring event with all its dynamic structure, this structure again being supported over the state of the brain taken as a wave, it is perhaps better to view this structure – also integral to specifying the current perception – as simultaneously supporting a "resonance" with all such similarly structured virtual events. How this virtual, past event enters the current brain state, gradually becoming an image, is the subject of a good deal of discussion in *Matter and Memory*, and I will acknowledge that it forms perhaps the most difficult aspect of Bergson's theory, but Bergson at least has a theory of the origin of the image of external world, and this is the memory-implication that must be dealt with.

We must be clear that we do not have here, in this time-extent of being, merely a long string or collection of separate "images." This was Sartre's (1962) mistake with respect to Bergson. Sartre seized on Bergson's referencing the universe as "an aggregate of images." He failed to realize that this is only one mode of speaking, referencing our normal mode of knowledge of the world in which to us, everything is necessarily an "image" – the "brain" is an image, a "neuron" is an image, a "fly" is an image, an "atom" – each but a partial aspect of the universal field. He entirely failed to grasp Bergson's deployment of another mode of viewing the universe – the holographic mode – as a vast interference pattern, a field non-image-able, wherein (unless by passing a reconstructive wave through) there are and can be no images. But clinging exclusively to this "aggregate of images" conception, Sartre was convinced that Bergson's memory consisted too of a vast set of separate images for which, he accused Bergson, there is no method of association.

In reality, we are talking about the extension in the past of the experience of a four-dimensional being. The two-week long canoe trip I took in Canada is this extended structure. Via picking up a paddle in the present, I may redintegrate from this structure an image of paddling down the lake. In the act of lighting a fire, I may redintegrate a similar fire-starting effort from that trip. All the "images" of this trip are in fact already integrally "associated," for they form part of one vast, indivisible structure of experience from which some aspects can be selected, seemingly separately and arbitrarily, via redintegration. Similarly, all the elements or "aspects" of coffee stirring – spoon, swirl, cup, clink, hand, cream, brown – are integrally associated as intrinsic parts of the invariance structure of this time-extended event: there must be an instrument providing the force, a liquid giving resistance, a swirl as a necessary consequence of the spoon's motion, etc. These components, contrary to the connectionist approach, do not need to be "associated" via training trial after training trial in a connectionist net – there is neither any "error" adjustment necessary nor, for that matter, any such linking effort that makes any sense (cf. Robbins, 2008). Bergson's model is a completely different approach – a non-associationist approach – to memory.

Hofstadter and Sander, one can suspect, intuitively knew what is actually required for the remindings underlying their constantly emerging analogies. Video recordings of all experience, etched upon neural structure, were implausible, but they knew they required the "storage" of that vast, detailed experience which in Bergson's framework – via the indivisible flow of time – is the natural extension of being. The redintegration mechanism forms at least the basis for retrievals from this mass of experience at various levels of abstraction. A precisely modulated reconstructive wave, i.e., a brain state supporting an invariance structure very close to an earlier event, may reconstruct that precise, original event – the time I was stirring soup over the campfire. A less precisely modulated state becomes a less precise wave sent through the entire mass and reconstructing all instances of stirring. We have what Gelernter (1994) conceived as a "stack" of events over which abstractions can be defined, or similarly, Galton's (1883) superposition of photographs of faces over which an abstract, generic face emerges. In this situation, the variants of each stirring event all wash out, the invariant remains – an abstract "stirring" that we term a "concept."

It is not in scope to explore the many aspects of memory this model brings up. Some of these aspects, for example the profound problem of explicit memory – the conscious localization of an event as an event in one's past – or the relation of such a model to connectionism, or its ties to other theories of the complete storage of events such as exemplar theory (e.g., Goldinger, 1998, 2007; Crowder, 1993), or a more detailed analysis of Bergson's concept of time and its relation to current physics, have been explored elsewhere (Robbins, 2009, 2014). The intent here has been to sharpen the problem of analogical reminding by showing that the memory mechanism needed to support it, based as it is in invariance defined only over flowing events, only exacerbates the storage dilemma. In this sharpening, however, I believe we are led naturally to the possibility that Bergson's model of experience and its "storage" may indeed be necessary.

References

Adelson, E., & Bergen, J. (1985). Spatiotemporal energy model of the perception of motion. *Journal of the Optical Society of America, 2*, 284-299.

Bar, M. (2007). The proactive brain: Using analogies and associations to generate predictions. *Trends in Cognitive Sciences, 11*, 280-289.

Bar, M. (2011). The proactive brain. In Bar, M. (Ed.), Predictions *in the Brain: Using the Past to Generate a Future*. Oxford: Oxford University Press.

Barsalou, L. W. (1983). Ad hoc categories. *Memory and Cognition, 11*, 211–227.

Barsalou, L. W. (1987). The instability of graded structure: Implications for the nature of concepts. In U. Neisser (Ed.), *Concepts and Conceptual Development* (pp. 101–140). Cambridge: Cambridge University Press.

Barsalou, L. W. (1993). Flexibility, structure and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A Collins, S. Gathercole, M. Conway, & P. Morris (Eds.), *Theories of Memory* (pp. 29-101). New Jersey: Erlbaum.

Barsalou, L.W., C.D. Wilson, & W. Hasenkamp (2010). On the vices of nominalization and the virtues of contextualizing. In B. Mesquita, L. Feldman-Barret, and E. Smith (Eds.), *The Mind in Context* (pp. *334-360).* New York: Guilford Press.

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

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

Casasanto, D. & Lupyan, G. (2015). All concepts are ad hoc concepts. In E. Margoulis & S. Lawrence (Eds.), The *Conceptual Mind: New Directions in the Study of Concepts* (pp. 543-566). Cambridge, MIT 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*, 185-211.

Clark, A. (2013). Whatever next? Predictive brains, situated agents and the future) of cognitive science. *Behavioral and Brain Sciences, 36*, 1-73.

Crowder, R. G. (1993). Systems and principles in memory theory: another critique of pure memory, in A Collins, S. Gathercole, M. Conway, P. Morris (Eds.), *Theories of Memory* (pp. 37-65). New Jersey: Erlbaum.

Dessoir, M. (1912), *Outlines of the history of psychology.* New York: MacMillan Co.

Dietrich, E. (2000). Analogy and conceptual change, or you can't step into the same mind twice. In E. Dietrich & A. B. Markman (Eds.), *Cognitive Dynamics: Conceptual and Representational Change in Humans and Machines* (pp. 265-294). New Jersey: Erlbaum.

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

Eich, J. (1985). Levels of processing, encoding specificity, elaboration, and CHARM. *Psychological Review. 92*, 1-38.

Elssaser, W. (1987). *Reflections on a theory of organisms*. Baltimore: John Hopkins University 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 A. Riegler, M. Peshl, & A. von Stein (Eds.), *Understanding Representation in the Cognitive Sciences* (pp. 158-163). New York: Plenum.

Freyd, J.J. (1987). Dynamic mental representations. *Psychological Review*, *94*, 427-438.

Galton, F. (1883). *Inquiries into Human Faculty and its Development.* London: Macmillan.

Gayler, R.W. (2003). Vector Symbolic Architectures answer Jackendoff's challenges for cognitive neuroscience. In Peter Slezak (Ed.), *ICCS/ASCS International Conference on Cognitive Science* (pp. 133-138). Sydney, Australia: University of New South Wales.

Gelernter, D. (1994). *The Muse in the Machine: Computerizing the Poetry of Human Thought.* New York: Free Press.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science, 7*, 155-70.

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.

Gick, M. & Holyhoak, K. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, *15*, 1-38.

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

Goldinger, S. (2007). A complementary-systems approach to abstract and episodic speech perception. In *Proceedings of the 16th International Congress of Phonetic Sciences* (pp. 49–54). Saarbruecken.

Hardcastle, V. G. (1995). *Locating Consciousness*. Philadelphia: John Benjamins.

Hofstadter, D. & Sander, E. (2013). *Surfaces and Essences: Analogy as the Fuel and Fire of Thinking*. New York: Basic Books.

Hohwy, J. (2013). *The Predictive Mind*. Oxford: Oxford University Press.

Hubel, D., & Wiesel, T. N. (1959). Receptive fields of single neurons in the cat's striate cortex. *Journal of Physiology*, *148*, 574-591.

Hubel, D., & Wiesel, T. N. (1978). Brain mechanisms in vision. *Scientific American, 241,* 150-162.

Hummel, J. E. & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Reviews, 12*, 487-519.

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

James, W. (1890). *Principles of Psychology.* New York: Holt and Co.

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, 1978.

Johnson, M. (1987). *The Body in the Mind: The Bodily Basis of Reason and Imagination.* Chicago: University of Chicago Press.

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

Kingma, I., van de Langenberg, R., & Beek, P. (2004). Which mechanical invariants are associated with the perception of length and heaviness on a nonvisible handheld rod? Testing the inertia tensor hypothesis. *Journal of Experimental Psychology: Human Perception and Performance, 30*, 346-354.

Klein, D. B. (1970). *A History of Scientific Psychology.* New York: Basic Books.

Kugler, P. & Turvey, M. (1987). *Information, Natural Law, and the Self-assembly of Rhythmic Movement*. Hillsdale, NJ: Erlbaum.

Lakoff, G. (1987). *Women, Fire, and Dangerous Things.* Chicago: University of Chicago Press.

Murdock, B.B. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, *89*, 609-626.

Mussati, C. L. (1924). Sui fenomeni stereocinetici. *Archivo Italiano di Psycologia, 3,* 105-120.

Nakayama, K. (1998). Vision fin de siΠcle: A reductionistic explanation of perception for the 21st century? In J. Hochberg (Ed.), *Perception and Cognition at Century's End* (pp. 307-331). New York: Academic Press.

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

Pribram, K. 1971. *Languages of the Brain.* New Jersey: Prentice-Hall.

Reichardt W, (1959). Autocorrelation and the central nervous system. In W. A. Rosenblith (Ed.) *Sensory Communication* (pp. 303-318). MIT Press Cambridge.

Robbins, S. E. (2002). Semantics, experience and time. *Cognitive Systems Research*, *3*, 301-337.

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

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

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

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

Robbins, S. E. (2013). Form, qualia and time: The hard problem reformed. *Mind and Matter, 2,* 1-25.

Robbins, S. E. (2014). *Collapsing the Singularity: Bergson, Gibson and the Mythologies of Artificial Intelligence*. Atlanta: CreateSpace.

Rogers, T., & McClelland, J. (2004). *Semantic Cognition: A Parallel Distributed Processing Approach.* Cambridge: MIT Press.

Rogers, T., & McClelland, J. (2008). Precis of: *Semantic Cognition: A Parallel Distributed Processing Approach*. *Behavioral and Brain Sciences*, *31*, 689-749

Sartre, J. (1962). *Imagination: A Psychological Critique*. (Translated by Forrest Williams). Ann Arbor, Michigan: University of Michigan Press.

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

Shaw, R. E., & McIntyre, M. (1974) The algoristic foundations of cognitive psychology. In, D. Palermo & W. Weimer (Eds.) *Cognition and the Symbolic Processes* (pp. 305-362). New Jersey: Lawrence Erlbaum Associates.

Sherry D., and Schacter, D. (1987). The Evolution of Multiple Memory Systems. *Psychological Review*, *94*, 439-454.

Spaulding, T, & Murphy, G. (1996). Effects of background knowledge on category construction. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *8*, 484-494.

Tulving, E. (1972). Episodic and Semantic Memory. In Tulving, E. & Donaldson, W. (Eds.) *Organization of Memory*. Academic Press.

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

Ullman, S. (1979a). *The Interpretation of Visual Motion*. Cambridge: MIT Press.

Ullman, S. (1979b). The interpretation of structure from motion. *Proceedings of the Royal Society of London, Series B*, *203*, 405-426.

Ullman, S. (1984). Maximizing rigidity: the incremental recovery of 3-D structure from rigid and non-rigid motion. *Perception, 13*, 255-274.

Ullman, S. (1986). Competence, performance and the rigidity assumption. *Perception, 15, 6*44-646.

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

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

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.

Watson, A. B. & Ahumada, A. J.. (1983). Model of human visual-motion sensing. *J. Opt. Soc. Am. A.*, *2*, 322-341.

Weiss, Y., Simoncelli, E., & Adelson, E. (2002). Motion illusions as optimal percepts. *Nature Neuroscience, 5*, 598-604.

Weiss, Y., & Adelson, E. (1998). Slow and smooth: A Bayesian theory for the combination of local motion signals in human vision. MIT A. I. Memo No. 1624.

Wheeler, M. (2008). Cognition in Context: Phenomenology, Situated Robotics and the Frame Problem. *International Journal of Philosophical Studies*, *16*, 323-49.

Zimmer, H. D., Helstrup, T., & Engelkamp, J. (2000). Pop-Out into Memory: A Retrieval Mechanism That is Enhanced with the Recall of Subject-Performed Tasks. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *26*, 658-670.