the mapping from phenomenal states to brains. To represent movement is not to have a moving representation – presumably we represent without using the very properties that are represented. But then why should we assume in this case that representations of a certain *type* of movement move in the very *way* that is represented? We are left with a pair of troubling facts: knowing the structure of one's phenomenal states does not seem to tell us enough detail about the structure of one's brain, and knowing the structure of a represented domain does not seem to entail anything about the brain either. Suppose *both* that Chasles' theorem governs phenomenal states, *and* that we were selected to represent objects in the world as obeying Chasles' theorem. What follows about brains?

Dynamics, not kinematics, is an adequate basis for perception

Andrew Wilson and Geoffrey P. Bingham

Department of Psychology, Indiana University, Bloomington, IN 47401.
{anwilson; gbingham}@indiana.edu

Abstract: Roger Shepard's description of an abstract representational space defined by landmark objects and kinematic transformations between them fails to successfully capture the essence of the perceptual tasks he expects of it, such as object recognition. Ultimately, objects are recognized in the context of events. The dynamic nature of events is what determines the perceived kinematic behavior, and it is at the level of dynamics that events can be classified as types.
[SHEPARD]

Roger **SHEPARD** has produced a fascinating account of how one might go about functionally representing the world. However, he has failed to successfully motivate his account. His evolutionary motivation is transformable into a story in which the advantage lies in the organism being able to flexibly perceive the world as presented to it (to support functionally effective action), rather than perceive the world represented by it. Also, it is not clear how conception should be uniquely related to perception and, therefore, how the study of conceptual problem solving can tell us anything about perception. SHEPARD uses his findings from studies on thinking to make claims about the nature of perception. We take issue with these claims.

SHEPARD notes, "My own searches for universal psychological principles for diverse perceptual-cognitive domains have been unified by the idea that invariance can be expected to emerge only when such principles are framed with respect to the appropriate *representational space* for each domain" (Introduction, pp. 581–82). His space is defined by canonical "landmarks" and the kinematic transformations that can transform them to match what is being perceived. We agree with his sentiment, but will argue that discrete kinematic transformations on time-independent landmarks do not make a suitable space. The space must contain continuous spatial-temporal forms because these (dynamically determined) "event forms" allow events to be identified as an example of a type, enabling recognition *of* the event and of objects *in* the event. A smooth gray surface could equally well be on a ball made of styrofoam or rubber. The continuous forms of motion exhibited when the objects are dropped on the floor allow them to be recognized for what they are.

Any motion, according to **SHEPARD**, is compatible with a dynamical account, given *arbitrary unseen* forces in nature. In his research where two discrete images are presented serially to subjects, they react in a manner consistent with a mental transformation of the initial object into the second via the kinematically optimal path. Not assuming the "arbitrary" forces makes kinematics the more empirically adequate account of these results. SHEPARD uses this as evidence that kinematics is the internalized transformation rule, and bolsters this claim with the additional claim that only kinematics, and not dynamics, is visually specified and, therefore, available for internalization. If dynamic properties, such as mass distribution and its consequent inertial properties, are not uniquely specified visually, they cannot serve as a reliable, internalizable set of transformation rules.

An error implicit in this argument is the suggestion that the forces in nature are arbitrary. They are not. Dynamicists study and describe reliable regularities in nature, configurations of force laws that produce the invariant forms of events that make the event recognizable. **SHEPARD**'s mistake is to exclude these consequent forms as "unseen." The effect of gravity on motion is perceptually salient, and observers are competent in using this information to identify events. The dynamics are immanent in the recognizable event whose spatial-temporal form is determined by those dynamics. That is not to say that gravity per se is necessarily recognized for its specific role in constraining the motion of a bouncing ball. Pre-Socratic philosophers need not have recognized gravity, but they certainly would have recognized bouncing balls, and observers fail to recognize bouncing balls as such when gravity has been altered (Twardy 1999; submitted). Events are recognizable according to the invariant kinematics, which, in turn, is determined by the underlying dynamics (Bingham 1995; 2000; Bingham et al. 1995; McConnell et al. 1998; Muchisky & Bingham, in press).

Two fundamental attributes of representational accounts are the ability of the representation to become causally de-coupled from that which it represents, and the low-dimensional nature of any computationally feasible representation. In **SHEPARD**'s account, these are achieved by reducing the visual space to a manifold of that space, defined by templates and geometric transformations for restoring something like the full dimensionality of our perceptual experience. In order to generate a veridical perceptual experience, the representation begins with its templates (and so must be in an appropriate initial position in the representational space), and transforms that template according to its geometrical rules to match the object being perceived. For this to work, both steps are highly constrained by the details of the event itself. If they are not, then the two attributes noted above mean that the veridicality of the perception is in doubt. In order for a representation to serve as the basis for tasks such as object identification it must be capable of generating constant matches and rematches between actual objects in motion and represented-object-transformed-by-geometrical-rules.

Imagine viewing a leaf blown about on a gusty day. By **SHEPARD**'s account, the perception of this leaf would require the computation and constant re-computation of the match of the leaf to a canonical "leaf." This computation would filter out the kinematic details specific to gravity and aerodynamics. Now imagine the perception of dozens of leaves on a gusty fall day. The representational structures would have to be generating constant, real-time updates for each of the large number of structurally similar objects all moving very differently in terms of the momentary directions and orientations (but moving in the same fashion in terms of the dynamically determined type of event). In Shepard's account, the computations would be defined kinematically, whereas what we actually see is best defined *dynamically*. This means that an unconstrained reconstruction could easily fail to use the appropriate kinematic mapping for the event's dynamic properties.

The reason for this can be thought of this way: kinematics is a description of a particular, local motion, one specific trajectory through a state space. Dynamics is the abstract description of the state space itself, describing the entire set of possible motions that correspond to a type of event. A pendulum can exhibit very different kinematic behavior. It can spin in circles, or swing in an arc. Both of these are captured in a single dynamical description, and it is only at the level of dynamics that these two motions can be classified as the same type of event, namely a pendulum event (Bingham 1995). Similarly, on a fall day, one does not merely see a collection of leaves; one sees a lot of falling wind-blown leaves (as compared to leaves that are merely falling on a calm day).

SHEPARD's representational model is clear and well formalized, and has made conceptual contributions to recent work in computer vision (Edelman 1999). However, the model is demonstrably not about human perception. Edelman's implementation is interesting for object recognition in image based processing. Human vision is not based on static images, however, and experiments implying kinematic reconstructions of potential transformations between discrete images are not a fair test of the human visual system's event perception capabilities. Using such experiments to separate the role of the perceiver from that of the dynamic world removes nearly all of the relational information normally available to, and utilized by, the perceiver in event perception, and changes the nature of the task. The world is such that the nature of the perceived event is specified adequately while the organism is causally coupled to that event. When not coupled to the event, things such as imagination, dreams, or imagery are, as SHEPARD claims, likely to be derivations based on the way the event was originally perceived. If kinematics is insufficient for the original perception, it is therefore unlikely to be sufficient for any related task.

Internalized constraints may function as an emulator

Margaret Wilson

Department of Psychology, North Dakota State University, Fargo, ND 58105.
margaret_wilson@ndsu.nodak.edu
http://www.melange.psych.ndsu.nodak.edu

Abstract: Kubovy and Epstein's main quarrel is with the concept of "internalization." I argue that they underestimate the aptness of this metaphor. In particular, an emulator which predicts unfolding events can be described as an internalization of external structures. Further, an emulator may use motoric as well as perceptual resources, which lends support to Hecht's proposal.
[HECHT; KUBOVY & EPSTEIN]

Has evolution caused our perceputal systems to "internalize" the constraints of the physical world? KUBOVY & EPSTEIN (K&E) answer this in the negative, but more on grounds of phrasing than on grounds of fact. To address their quarrel with this question, it may be useful to divide it into three questions:

(1) What abilities or tendencies have been built into our perceptual machinery? Do they mirror the constraints of the physical world?

(2) Are these tendencies represented explicitly in the form of rules, or are they present implicitly in the functioning of the perceptual machinery?

(3) Do these tendencies deserve to be called internalized constraints? Is any explanatory power added by conceptualizing things in this way?

K&E appear to agree with SHEPARD on the first question. They insist that constraints are something out in the world, not in the head; but they also acknowledge that to benefit from these constraints the perceptual system must have been shaped to produce output concordant with the constraints.

It is with the second question that the quarrel begins. K&E read SHEPARD as favoring the explicit view, with "mental contents actively engaged in the perceptual process." In contrast, they themselves prefer an account in which the system does not *follow* rules, but rather *instantiates* them.

How much, though, really hangs on this question? It is surely an interesting question in and of itself (though a notoriously difficult one to answer with any satisfactory clarity); but does the answer have substantial consequences for SHEPARD's idea of internalization? K&E concede that "the visual system proceeds *as if* it possessed knowledge of kinematic geometry." But if there is fundamental agreement on the behavior of the visual system, and on how closely it tracks the properties of the physical world, it seems to me that Shepard's point is carried. The precise nature of the machinery which does this tracking can be left to future investigation. (The importance assigned to this question, though, may come down to personal preference. Some investigators, such as myself, have considerable tolerance for the "as if" style of explanation so pervasive in cognitive psychology, while others find it unacceptable.)

The quarrel continues with the third question. K&E object to the word "internalization," and suggest that it is an appealing metaphor but ultimately not an appropriate one. Yet, while they suggest a number of reasons why such theoretical terms *can* be problematic, they don't clearly state an objection to this particular term. I suspect that their objection in fact rests on their answer to the second question. The "internalization" phraseology suggests that the constraints of the external world are "things" that have been lodged in the mind in the form of explicitly stated rules.

But the internalization metaphor need not be read this way. External constraints – the laws of physics and so on – are not themselves explicitly represented rules. Instead, they are implicitly instantiated in how the system, the physical world, behaves. And in response to evolutionary pressure, based on the realities of living in such a world, a parallel set of "constraints" has sprung up within the cognitive system, again instantiated in how the system behaves. In such a situation, where external constraints cause the emergence of a mirroring set of internal "constraints," it seems to me that "internalization" is a particularly apt metaphor. (K&E make reference to the work of Lakoff & Johnson 1990 regarding the metaphorical basis of abstract concepts, but neglect the point that these metaphors are so pervasive, systematic, and enduring because they often track reality so well.)

Is any useful purposed served, though, by this "internalization" terminology? I suggest that there is, not least because it invites contact with a related set of ideas gaining currency in the field. A problem faced by any physical system that must interact with the world in real time, be it a remote-controlled factory robot or a human body, is that of making corrections and adjustments in response to feedback that is delayed by the time required for signal transmission. Even a slight delay in feedback can result in overcompensation for errors, which then require further compensation, and so on. One possible engineering solution is the use of an "emulator," a mechanism within the control system that mimics the behavior of the situation being acted upon, taking afferent copies of motor commands and producing predictions of what should happen (e.g., Clark & Grush 1999; Grush 1995).

To be useful, an emulator of course needs to be a successful mimic of the external situation. And to be a successful mimic, it almost certainly needs to be structurally isomorphic, for certain relevant properties, to the situation being emulated. Although much that happens in the world is unpredictable, there is also a good deal of regularity and redundancy that could be exploited by an emulator. It does not seem far fetched, in such circumstances, to say that the structure of the situation has been internalized. This may be just what has occurred in the evolution of the human visual system.

If the human visual system does indeed use this kind of an emulator, then the percepts we experience under conditions of ambiguous or "gappy" input would presumably reflect the predictions produced by the emulator. This may help to bridge the gulf between the minimalist stimuli of the laboratory and the more robust input usually provided by ordinary perception. The emulator functions in both situations, but is only allowed to truly determine the content of the percept when there is a temporary absence of reliable input.

Further, it is possible that this same emulator, run off-line, is responsible for the phenomenology of mental imagery (cf. Grush 1995). If this is true, and if there is indeed an isomorphism between emulator and world on some restricted set of relevant properties, then "mental rotation" may in fact be a mental process that is isomorphic to the actual rotation of a physical object.