Regularities of the physical world and the absence of their internalization

Heiko Hecht
Man-Vehicle Lab, Massachusetts Institute of Technology, Cambridge, MA 02139
hecht@mit.edu **www.uni-bielefeld.de/zif/heiko.html**

Abstract: The notion of internalization put forth by Roger Shepard continues to be appealing and challenging. He suggests that we have internalized, during our evolutionary development, environmental regularities, or constraints. Internalization solves one of the hardest problems of perceptual psychology: the underspecification problem. That is the problem of how well-defined perceptual experience is generated from the often ambiguous and incomplete sensory stimulation. Yet, the notion of internalization creates new problems that may outweigh the solution of the underspecification problem. To support this claim, I first examine the concept of internalization, breaking it down into several distinct interpretations. These range from well-resolved dynamic regularities to ill-resolved statistical regularities. As a function of the interpretation the researcher selects, an empirical test of the internalization hypothesis may be straightforward or it may become virtually impossible. I then attempt to cover the range of interpretations by drawing on examples from different domains of visual event perception. Unfortunately, the experimental tests regarding most candidate regularities, such as gravitational acceleration, fail to support the concept of internalization. This suggests that narrow interpretations of the concept should be given up in favor of more abstract interpretations. However, the latter are not easily amenable to empirical testing. There is nonetheless a way to test these abstract interpretations by contrasting internalization with the opposite concept: externalization of body dynamics. I summarize evidence for such a projection of body constraints onto external objects. Based on the combined evidence of well-resolved and ill-resolved regularities, the value of the notion of internalization has to be reassessed.

Keywords: event perception; evolution; internalization

Introduction

Shepard's (1994) claim that our minds reflect the very same principles that govern the universe is appealing indeed. According to this claim, the mind has internalized universal principles (regularities) that allow it to disambiguate situations that would otherwise be unsolvable. Provided the world is not changing, such universal principles are very efficient. For the visual system, this explains why we can make sense of stimuli that by themselves do not suffice to specify our perceptions. As I will show in this paper, as appealing as this claim is, it has two interpretations that need to be distinguished. Both become problematic when subjected to closer scrutiny.

There is a troubling duality to Shepard's internalization hypothesis. On the one hand, the convincing example of an inner circadian rhythm suggests he takes internalization to mean that a well-defined physical regularity is also independently present in the organism and allows behavior consistent with the regularity even if it is no longer there (as is day or night for people in a dark cave). This example is quite unique, and other examples, such as kinematic geometry, do not assume any exact mirroring of a physical law in the perceptual or behavioral outcome. To the contrary, kinematic geometry supposes a good deal of abstraction from movements that are found in the physical world. The two examples are symptomatic for two vastly different readings of the internalization hypothesis. The former I call the literal interpretation. The latter I call the abstraction interpretation.

For the literal interpretation, to determine whether or not some principle or regularity of the physical world has been internalized, three things have to be true: (1) First, there has to be a regularity in the world that can be assessed independently of our perceptions. (2) Second, our behavior and/or percept has to be compatible with this regularity as established by empirical observation. (3) Third, additional evidence is needed to show that the percept has come about by virtue of internalization and not by some other learning process. The first two steps are comparatively straightforward while the third is very tricky. Fortunately, it does not have to be resolved when approaching the problem from a falsificationist point of view. As long as one and

Heiko Hecht is a post-doctoral associate working in the Man-Vehicle Laboratory at the Massachusetts Institute of Technology, Cambridge, MA. He has published on issues of motion perception and dynamic event perception. From 1996–1997 he was scientific coordinator of an international project on Evolutionarily Internalized Regularities at the Center for Interdisciplinary Research in Bielefeld, Germany.

  0140-525X/01 $12.50

two are the case, the (literal) internalization hypothesis survives the test.

A less literal reading of the internalization hypothesis causes much more trouble. It is also what Shepard most likely had in mind. To arrive at perceptual-cognitive universals (Shepard 1994), a mere copy of an often seen movement or event is clearly insufficient. Internalization as a process of abstraction is geared toward finding default solutions. Whenever the stimulus is ambiguous or ill-defined, as in apparent motion, an internalized default influences the percept. In this evolutionary process, geometry has been more deeply internalized than physics (Shepard 1994). Thus, an almost paradoxical relationship between the degree of internalization and its palpability is postulated. The deeper an invariant is internalized the more abstract it has to be. The process of internalization is then not just an inductive process, but also, by definition, one where the best examples are the least well-defined.

Unfortunately, this abstraction interpretation virtually annihilates the three requirements that hold for the literal interpretation: (1) The physical regularity no longer needs to be crisp and stateable as a physical law. (2) It becomes much harder to state what empirical behavior actually contradicts an abstract internalized rule. Together, this leads to potential research problems that are addressed below. (3) With the broader interpretation of the internalization hypothesis, the issue of internalization versus learning may reach beyond what can be tested empirically (see Schwartz, this issue).

An empirical assessment of Shepard's internalization hypothesis – and this is my quest – is thus inextricably tied to its interpretation. The literal and the abstract interpretation can be taken to represent the two ends of a continuum, within which Shepard is hard to place. Because of this difficulty, I resort to the strategy of evaluating a variety of internalization candidates that range from very narrow to very broad readings of internalization. For all of them I stick to the general premise that the internalized knowledge comes into play when the percept is ill-defined or when conflicting cues have to be resolved. Since a hypothesis should not be tested with the examples that were drawn up for its initial support, I pick some domains that have not been considered by Shepard but in my opinion constitute good cases for potential internalization. These are some literal laws of physics as well as some examples from the more abstract domain of intuitive physics. I hope thus to analyze Shepard's claims and to elucidate the concept of internalization.

1. Classifying candidates for internalized regularities that govern the physical world

A natural strategy for an empirical test of *the* internalization hypothesis would be to first examine different types of regularities at differing levels of abstraction, and then to test whether our percepts reflect these regularities. To evaluate a given candidate regularity, three questions should be answered. First, to what degree does it describe the physical world, that is, are there exceptions or is it universally true? Second, what is its level of complexity? A very complex natural law may hold without exception but it might be impenetrable to the visual system and appear inconsistent. Third, what is the degree of abstraction that is involved in a given internalization hypothesis? Possible candidates can be grouped as a function of how they score on these questions. I distinguish the following groups: potentially internal regularities that are close to the laws of physics, such as dynamic invariants; specific but highly abstracted rules, such as kinematic geometry; more general rules, such as the Gestalt principles; and unspecific and highly abstracted regularities, such as Bayesian probabilities.

1.1. Dynamic regularities

The strongest case for the internalization of physical regularities would be made if a simple invariant that holds in the physical world guides our perception. The fact that light usually comes from above seems to fit this category perfectly. There are many examples of unexpected and unnoticed artificial illumination from below, however, that can perceptually invert the scene. Valleys are turned into mountains and vice versa (Metzger 1975; Ramachandran 1988a). Light does not always come from above, though *generally* it does and this may have prompted the visual system to use that assumption when the stimulus is not very rich, as when looking at photographs or masks of human faces. This "illusion" has not been reported in more ecological settings, but that poses no threat to the internalization idea. At this literal level, only the failure to recur to plausible regularity assumptions would pose a threat to the hypothesis. I contend that such data is there to be used and can be gleaned from studies of intuitive physics (see the section on candidates for internalization). Unfortunately, the other good candidates, such as the constant gravitational acceleration of falling bodies, do not seem to support the notion of successful internalization.

1.2. Geometric regularities

The most detailed internalization hypothesis that Shepard (1984) has put forth is that of kinematic geometry. He proposed the internalization of geometric principles pertaining to group theory at a high level of abstraction. These principles, which prescribe, among others, circularly curved motion paths, are thought to act as a general default that influences perceiving, imagining, and dreaming. This abstraction variant of the internalization hypothesis remains very controversial (see Todorovič, this issue). The empirical evidence gathered by Shepard himself (e.g., Lakatos & Shepard 1997) causes confusion about what exactly is meant to be internalized. Three different views are possible and leave a number of back doors open to maintain abstract internalization: (1) The crisp law of geodesic movements could have been "imperfectly" internalized. (2) A general, imperfect law could have been perfectly internalized. (3) A fuzzy general law could have been imperfectly internalized. It is not hard to see that empirical data can be imperfect in multiple ways and still be compatible with Shepard's proposition. This issue will be taken up in the section on kinematic geometry.

Another example of how the visual system exploits knowledge about geometric regularities has been put forth by Bingham (1993). He found that observers use the shape of unfamiliar trees to judge their absolute size. To do so general relationships such as ratio of trunk to branch size, number of branches, and so on are exploited. Here shape can even override horizon-ratio information, which is normally very informative (Rogers 1995), at least as far as pictures are concerned. Thus, there is evidence outside the realm of kinematic geometry demonstrating that observers make use of prior shape knowledge to judge absolute size. However, this potential role of geometric regularities in perception does not entail that the geometric knowledge is internalized (vs. learned), or universal.

1.3. Gestalt principles

A next step of abstraction is reached when Gestalt principles are interpreted as internalized regularities. Gestalt psychology shares the conviction that the different Gestalt principles reflect very general regularities. The well-known Gestalt principles, such as grouping by similarity, by common fate, or, more recently, by uniform connectedness (Palmer & Rock 1994) will not be discussed here because traditionally they have not touched on the issue of internalization, probably because Gestalt psychologists were more concerned with a static description of phenomena and physiology than with evolutionary processes. Some Gestalt theorists (e.g., Metzger 1975) have even treated Gestalt principles as the very conditions that make perceptual psychology possible. Thus, Gestalt psychologists acknowledge the pervasiveness and *a priori* nature of a whole list of principles and they do not single out one, such as kinematic geometry. They also do not stress the processes of internalization but conceive of them in an almost Kantian fashion as preconditions of experience. The notion of an internalization could be taken as an evolutionary explanation of the origin of general principles, including Gestalt principles.

1.4. Statistical regularities as a special case of maximal abstraction

The most abstract way to describe regularities of the physical world that are reflected in the visual system consists in pointing out mere statistical relationships. Typically, if an equal distribution assumption is made, we can predict which views of objects are likely and which ones are rare. For example, it is extremely unlikely that we see a pencil exactly head-on such that it produces the retinal image of a circular patch (provided by monocular viewing). Consequently, a circular retinal patch is normally not interpreted as a pencil but rather as a round object. The notion that the visual system "knows" generic views from accidental ones has been put forth in Bayesian approaches to perception (e.g., Albert & Hoffman 1995; Hoffman 1998), which postulate that the organism makes use of prior information about the world. For instance, Hoffman (1998) describes such knowledge as a list of rules that the visual system applies to the stimulus, such as "interpret[ing] a straight line in an image as straight line in 3-D."

This reconstructionist view gathers support from Shepard (1987b), who suggests how such prior knowledge could have developed by a process of internalization. His explanation draws on probabilistic aspects of nature and processes of stimulus generalization within the organism. The likelihood of responding to a new stimulus the same way as to a different previously learned stimulus (generalization) depends on the proximity between the two stimuli in psychological space. According to Shepard, this function is not equivalent with discriminability but reflects the anticipated consequences of the reaction toward the stimulus class. The function is exponential and supposedly reflects a universal law that is as ubiquitous for animate beings as Newton's law of gravitation is for inanimate objects. This is a good argument for why internalized laws are poorly resolved and may have to be imprecise. Unfortunately, it makes it very hard to interpret empirical data that do not quite fit the supposed regularity. On the one hand, such data could be taken to mean that a well-resolved regularity has been internalized poorly. On the other hand, it could mean that the regularity has been abstracted and then internalized perfectly. Without a set of independently derived abstraction rules, we cannot favor one interpretation. And since such rules have not been formulated, the internalizationist's foregone conclusion is that internalization has been demonstrated. This seems to prompt Proffitt and Kaiser (1998) to conclude that the visual system has not internalized (well resolved) dynamic constraints but rather (coarser) geometric concepts.

It is easy to arrive at this conclusion under the premise that internalization has to be perfect, but this is most likely not the case. If this be demanded, existing empirical evidence suffices to falsify claims of internalization of both dynamic and geometric concepts. To support this point, I will summarize representative empirical evidence showing that our percepts are often only approximated by such concepts. Neither dynamic invariants (gravity, horizontality) nor optical invariants (tau) nor geometric rules (geodesic paths) predict our perceptions with satisfactory accuracy. Shepard has tried to turn this vice into a virtue by introducing a process of abstraction into the concept.

2. Three example cases for internalization

If universal but specific regularities can be found which appear to guide our perception whenever underspecified, a case could be made for internal knowledge and maybe even for a process of internalization. Once this is done, more abstract, generalized versions of physical regularities can be considered. Thus, I first examine gravity and horizontality as potential dynamic regularities. Based on the negative results, the would-be universality of apparent motion trajectories will then be reconsidered, reevaluating the example of kinematic geometry.

2.1. Gravity as a cue to absolute size and distance

The force of gravity is not only ubiquitous but also accelerates all terrestrial objects at a constant rate. Gravity is thus a prime candidate for a specific constraint that the visual system might have internalized to disambiguate perception. The internal knowledge in this case would be indirect. If observers judge the absolute size of an object more accurately when they see it fall, they may use implicit knowledge about gravitational acceleration to perform this task. Saxberg (1987) and Watson et al. (1992) suggested that observers do in fact estimate the absolute size and/or distance of objects by relying on the monocular cue of gravitational acceleration as is present in projectiles in flight, pendulum motion, fluid wave motion, and others. Saxberg, for example, showed that one could estimate the absolute distance to an object from four retinal image variables: the vertical and horizontal components of the object's retinal velocity and the vertical and horizontal components of its retinal acceleration. The estimation of absolute distance is even simpler when the object's motion is vertical; the vertical retinal acceleration in that case is proportional to the gravitational constant (Watson et al. 1992). Thus, one can in principle estimate the absolute distance to a freely-falling object. This estimate becomes more complex when friction plays a role as is the case for light and fast objects. However, for most inanimate objects within our space of action, air resistance has comparatively small effects.

Empirical results, however, showed that observers do not behave as if they make use of some knowledge about gravity (Hecht et al. 1996). Computer-simulated events of free falling objects revealed that observers were not very good at scaling absolute size and/or distance. Balls of different diameters and at different distances from the observer were simulated to rise, climb to their apex position, and then fall back down. Two categories of events were used, accelerating balls and constant-velocity balls. The latter had the same event-durations and average velocities as the former; however, only accelerating stimuli could be used to scale the distance of the event. Figure 1 depicts the different position/time diagrams for a subset of the stimuli whose distance observers had to judge. Observers did not perform better on the accelerating trials than on constant velocity trials, but both were considerably better than static versions of the stimuli. It can thus be ruled out that observers used projected size as a cue, which is always correlated negatively with simulated distance and positively correlated with simulated size.

Thus, observers do not utilize specific knowledge about gravitational acceleration to a sufficient degree. It remains possible that some abstraction of this regularity has nonetheless been internalized. Average image velocity is necessarily and negatively correlated with simulated distance whenever the apex point is shown. The fact that size and distance judgments in the constant velocity condition were significantly better than chance shows that observers can, in fact, make use of the average velocity cue. They behave as if they were abiding by a simple heuristic such as "objects that produce fast retinal motion are relatively close to me."

The hypothesis that gravitational acceleration has nonetheless been internalized could be salvaged by assuming that, for some reasons, fast-moving objects for whom air resistance is no longer negligible, have determined the internalization process. In this case, air-resistance is sometimes considerable and the effects of gravity vary depending on density and size of the falling objects. Drag, for instance, increases geometrically with object velocity. For a baseball moving at 80 miles/hr, the drag is about 70% of the ball's weight (Brancazio 1985). Thus, the visual system, instead of having to adjust for drag, might have adopted a cruder mechanism reflecting the fact that moving/falling objects give rise to higher retinal image velocities at closer distances. This relationship usually holds no matter how the object moves and whether the object is accelerating or moving at constant velocity. Presumably, observers are sensitive to this fundamental relationship and the visual system could use this abstract information to disambiguate percepts of distance.

In sum, we have to reject the falsifiable hypothesis that observers have internalized detailed knowledge about the rate of gravitational acceleration. The less specific case is still possible, but it may also be immune to criticism (see Fig. 4).

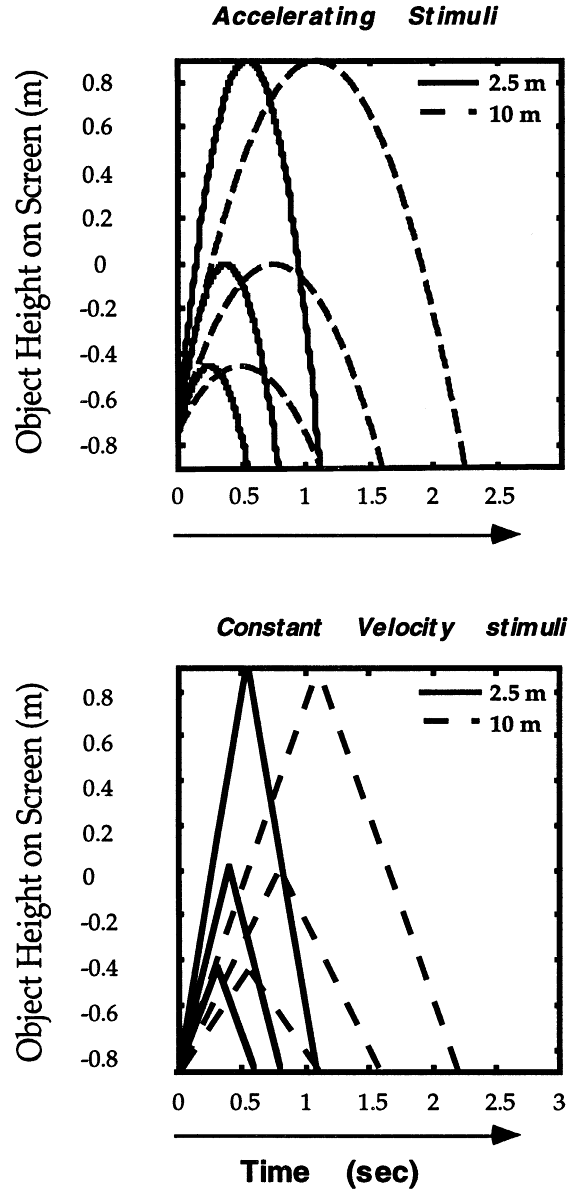

Figure 1. Projected ball trajectories for accelerating and constant velocity trials. Vertical position on the display screen is plotted as a function of time. The left panel shows the trajectories for the accelerating condition for simulated distances of 2.5 and 10 m. There were three different apices in this experiment, one at the top of the screen, one three-quarters of the way to the top, and one half of the way. The right panel shows the trajectories for the constant velocity condition for simulated distances of 2.5 and 10 m (adapted from Hecht et al. 1996, p. 1070).

2.2. The law of horizontality and the water-level task

When shown a tilted container people often fail to appreciate that the surface of the contained liquid should remain horizontal with respect to the ground. A typical example of the paper-and-pencil version of this Piagetian task is depicted in Figure 2. Subjects are asked to draw in the surface of the water such that it touches the dot on the right side of the container. About 40% of the adult population draw water-levels that deviate by more then 5% (α) from horizontal (for an overview see Liben 1991). The failure to solve this water-level task correctly is quite robust

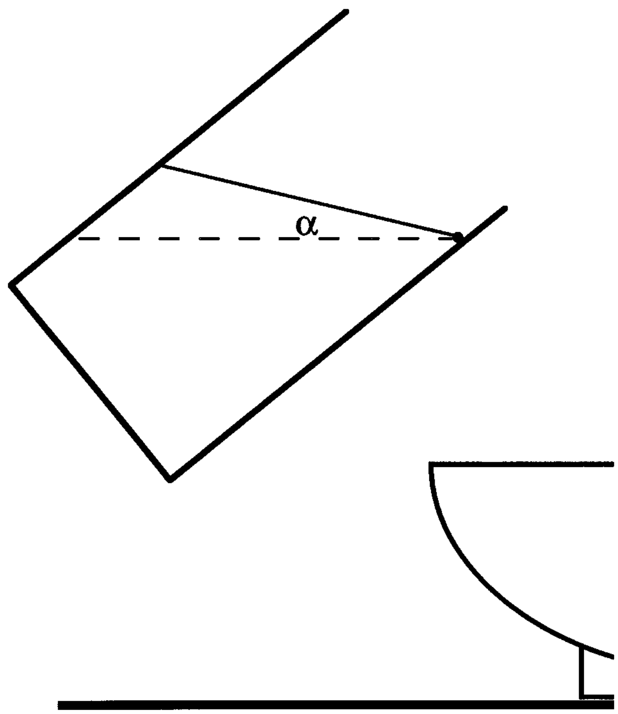

Figure 2. The water-level task. Observers have to draw the surface of the water. They are told that the beaker is at rest and filled by as much water as needed to make the surface touch the dot on the right of the container. The dotted line indicates the correct solution, which was only produced by one-half of all subjects. The solid line depicts a typical answer given by the other subjects.

across presentation contexts and does not appear to be an artifact of the technique that is chosen to communicate the task.

Howard (1978) presented apparent motion sequences of photographs depicting horizontal and oblique water-levels and asked subjects to report whether the sequence represented a natural or an unnatural event. Using an animated version of the task did not improve performance (Howard 1978; McAfee & Proffitt 1991). The would-be internalization of horizontality did not come to the fore when pouring events with impossible water-levels were shown by virtue of tilting the camera when the scenes were videotaped. Thus, a variety of methods used to assess the explicit and implicit knowledge of the horizontality invariant produced the same results: the regularity that liquid surfaces at rest remain invariably horizontal with respect to the ground cannot be taken to be internalized by our visual system.

If the horizontality of liquids has been internalized in a more ephemeral manner, visual experience may be required before the internalized regularity manifests itself in behavioral data. Thus, one might argue that with sufficient experience the "illusion" should disappear. However, the opposite is the case. Experienced waitresses and bartenders reveal stronger biases than the average population; they accept water levels as natural that deviate even more from horizontality (Hecht & Proffitt 1995; but see Vasta et al. 1997).

In conclusion, two examples of physical regularities have failed to influence perceptual judgments. The empirical data have thereby failed to fulfill a precondition for the possible internalization of theses regularities. Percepts were vastly different from the defaults that hold in the environment with few exceptions. Knowledge about gravity, per se, is not used to scale absolute size and distance of objects. However, the general negative relationship between retinal velocity and distance of a moving object could still be said to be internalized. A much stronger case against the internalization hypothesis is represented by the water-level literature. The regularity that liquid surfaces remain invariably horizontal when at rest is as consistent as the diurnal cycles, it has no exceptions. Nonetheless, a substantial proportion of observers misjudge water-levels, indicating that they have by no means internalized this particular regularity.

Thus, the failure to exploit these rather concrete regularities contradicts the literal reading of Shepard's hypothesis. However, neither free-fall nor the water-level task have a geometric solution that differs from the laws of classical mechanics. Thus, if one excludes these examples from the domain of internalization theory, the latter may not be threatened. Such an exclusion might be put forth on grounds of insufficient underspecification of the percept in the case of gravity to scale objects and on the importance of cognitive factors in the water-level case. This would rescue the literal interpretation but seriously reduce the scope of the hypothesis.

2.3. Kinematic geometry?

Let us now look at the abstraction reading of internalization. Given the negative results that were obtained when testing the comparatively simple regularities of gravity and horizontality, we search for evidence of internalization in situations where the percept is severely under-specified but not arbitrary. Such cases are very hard to find. Dreaming and imagining – although suggested by Shepard (1984) – appear to lack specification altogether and are, in addition, very hard to measure. Apparent motion seems to be the only appropriate case that could provide evidence for internalization of abstracted regularities.

For moving extended objects, Shepard (1984; 1994) claims that the perceived trajectory of an object's motion is determined by the geometrically simple geodesics. The model is based on the idea that a group of single rotations (SO[3]) can define the space of all possible three-dimensional (3-D) orientation differences (Carlton & Shepard 1990a; 1990b; Foster 1975b). Within this space, Chasles' theorem describes the simplest single rotation as follows: for any object displacement and orientation change, there exists one axis in space about which the object can be rotated, such that its initial position will be mapped into its final position. This helical motion in 3-D reduces to a single rotation (without concomitant translation) in the 2-D case.

However, the empirical evidence, including some of Shepard's own studies (McBeath & Shepard 1989), does not always support the geodesic model. McBeath and Shepard's data fell somewhere between the straight line path suggested by principles of energy minimization and the postulated geodesic path. Empirical apparent motion trajectories in 3-D especially often deviate considerably from the geodesic solution. Depending on the circumstances, perceived paths can be much closer to a straight line than to a geodesic curve even when the 3-D orientation of the motion plane necessary to specify the geodesic is properly judged (Hecht & Proffitt 1991). These results hold, of course, only when inter-stimulus intervals are sufficiently long so that geodesics could in principle be observed. A general model not based on kinematic geometry that could explain many of these results has been proposed by Caelli et al. (1993), who suggest a complex constraint-satisfaction procedure.

Recently, Shepard (1994) accommodated all deviating results in the apparent motion domain into his theory by claiming that we do not necessarily perceive motion in accordance with kinematic geometry whenever the percept is under-specified, as in apparent motion or imagery. Rather, he makes the weaker claim that kinematic geometry is "more deeply internalized than physics" (p. 7). This claim is too weak to be the basis for any predictions. If we take another intuitive physics example, what would Shepard predict for the following case? Imagine a marble that is rolled through a C-shaped tube, which is positioned horizontally on a table top. What will its movement path look like after it exits the tube? If the observer has internalized an approximation to Newtonian mechanics, she should imagine the marble to continue its path in a straight line perpendicular to the tube's opening (Fig. 3, case A). If on the other hand, curved geodesics are internalized, a curvilinear path might be preferred (Fig. 3, case B).[1] Empirically, many subjects erroneously think that the marble should continue to curve, presumably because it has acquired a curvilinear impetus (McCloskey et al. 1980). However, observers who make erroneous predictions prefer the correct straight path when confronted with visual animations of a variety of straight and curved trajectories (Kaiser et al. 1985a; 1992). Thus, only with less visual support are curved paths preferred. Have curved trajectories beyond Chasles' theorem been internalized, or does internalization fail here because it can predict all interesting outcomes?

3. Doubts about the epistemological status of internalization

The above examples show that the internalization hypothesis is in trouble. Taken together, those candidates of the internalization hypothesis that are amenable to empirical testing call for a revision of this concept. The literal interpretation of internalization is faced with heavy counter-evidence. The more likely abstract interpretation suffers from two very different problems that have to do with fundamental limits to its empirical verification. The first problem concerns the resolution or generality of the internalized rules, and the second concerns the need for a criterion that determines when internalization has occurred.

3.1. The resolution problem

If we say that an organism has internalized a particular regularity or rule, such as the periodicity of the circadian rhythm, we could refer to a very coarse level of resolution: some vague expectancy of day following night. On the other hand, we could mean that the organism has an internal clock and knows down to the minute when the sun will rise. The higher the level of resolution, the easier an empirical test. The level of resolution that we apply to the internalization hypothesis determines to what extent it is amenable to empirical testing.

Kubovy and Epstein (in this issue, p. 621) claim that the internalization hypothesis "has no obvious empirical content and cannot be tested experimentally." This is only true in its broadest reading. In support of Shepard, I not only hold that there are other readings that can be tested experimentally, I also claim that the more fine-grained the operationalization of his hypothesis, the easier it is to refute. For example, the hypothesis that we have internalized the rule that water surfaces at rest are always horizontal is a strong case that allows distinct predictions: we should resolve ambiguous perceptual situations in this context with errors toward a preference for horizontal orientation. On the other hand, the hypothesis that we have internalized some abstraction of this regularity would not necessarily put us in a position to use a few empirical observations of non-horizontal solutions as evidence against the internalization hypothesis. It may not be falsifiable at all if we cannot think of any behavior that could contradict the claim (see Popper 1935).

In Figure 4, I have tried to depict the relationship between postulated internalizations and hypothesis testing. The resolution of a given internalization hypothesis tends to correlate highly with its amenability to empirical testing. Highly resolved statements that claim generality are easy to

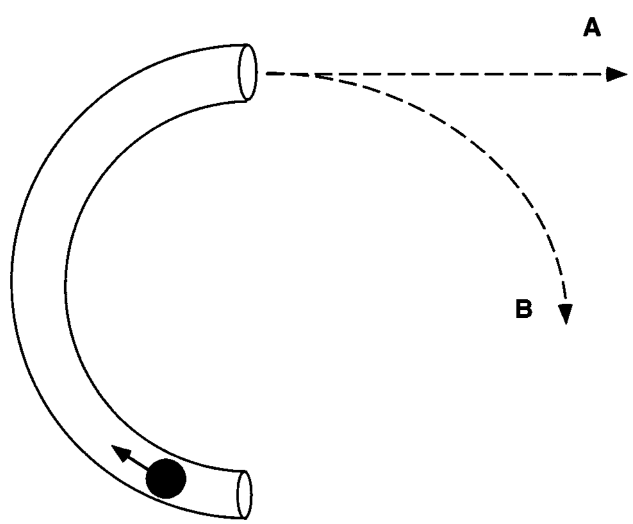

Figure 3. When asked to predict the path taken by a ball rolled through a C-shaped tube, many subjects mistakenly chose a curved trajectory (b) over the correct straight path (a).

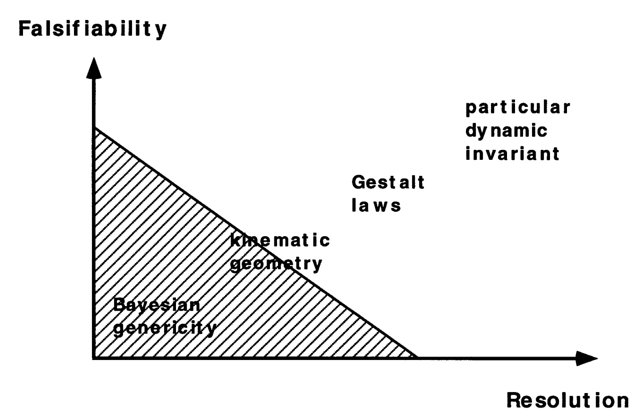

Figure 4. The internalization concept can be analyzed at different levels of resolution. The more specific and the better resolved an instantiation of the claim, the greater its chance to be found false after empirical testing. Ill-resolved claims that are hard or even impossible to falsify are immune to criticism (shaded area).

falsify and therefore desirable, as for instance the hypothesis that "all perceived apparent motion trajectories follow geodesic paths." Unfortunately, Shepard's internalization hypothesis is most appealing where it is least resolved. It may in fact be so appealing because it is immune to empirical testing. Also the notion of kinematic geometry and some Gestalt laws are on the brink of immunity as long as they are not supplemented with precise predictions, as for instance the Gestalt law of proximity ("objects in close spatial proximity tend to be perceptually grouped together"). Thus, when discussing *the* internalization hypothesis, we always have to add at what level of resolution we are making our argument. The above distinction of a literal and an abstract interpretation was an attempt to do so.

3.2. The criterion problem

The criterion problem refers to the content of the internalized knowledge. What type of knowledge can in principle be internalized? To answer this question, we need to narrow the concept of internalization. In its ill-resolved form, internalization can accommodate such diverse approaches as indirect and direct theories of visual perception. According to the former, without further assumptions, the visual system could not arrive at unique interpretations of the necessarily ambiguous retinal stimulus (Rock 1983; von Helmholtz 1894). Visual perception has to solve ill-posed problems that have no unique solution (Poggio 1990). Assumptions that transform ill-posed problems into well-posed ones are typically not arbitrary, and the percepts they create are not qualified by a question mark (as in the case of a Necker cube, whose percept can change momentarily), but are usually stable and distinct. In other words, the decoding of the stimulus information requires methods of induction (Braunstein 1994) and additional assumptions about the world, which the visual system has – in some broad sense – internalized (Shepard 1984). If this is the case, perceptual problems should not only be solved by the visual system, they should be solved in a manner consistent with laws that govern the physical world. Direct theories of perception (Gibson 1979) would phrase the same basic story rather differently. The makeup of the visual system, as developed through evolution, prepares it to pick up information relevant for proper action. In a sense, internalization is implicit here.

To sharpen the criterion for internalization, a minimal requirement seems to be that the organism must have had a chance to fail to internalize the knowledge in question. Truly universal *a prioris* of perception would thus not be candidates for internalization. Take, for example, the law of noncontradiction: if an object could at the same time exist and not exist, neither object recognition nor epistemology could work. Proponents of evolutionary epistemology argue that our evolutionary world knowledge has no choice but to work with these necessary constraints. This entails that they are also reflected in perceptual processes (Wächtershäuser 1987). This holds not only for laws of logic, but also for basic structural symmetries between the world and perception that may not be coincidental. Campbell (1987), for instance, points out the striking coincidence that almost all objects that reflect or absorb light also block our locomotion and, likewise, all objects that are permeable to light do not obstruct our locomotion. We can see and walk through air, to a lesser degree through water, and not at all through solid objects. This corresponds to the two fundamental constituents of the terrestrial environment, which Gibson (1979) construed to be media and surfaces (of substances). Their existence is too basic to be called "internalized" in any meaningful fashion. Likewise, the optics of the lens, the location of our eyeballs, and so on, impose constraints onto the visual system that need to be considered to understand vision, but do not qualify as examples of internalization.

Internalization also does not need to be an explicit or declarative knowledge structure. It can nontrivially be achieved by virtue of the makeup of the system. This has been conceptualized by direct perceptionists with the use of an analogy. The visual system acts like a smart device (Runeson 1977), as does, for instance, a speedometer. A speedometer does not measure time or distance, thus has no *knowledge* about speed, yet nonetheless "measures" speed by means of an induction current caused by the revolutions of the wheel and translated into the position of the speedometer needle. Taking advantage of induction is, in a manner of speaking, evidence for the fact that some principles of physics have been internalized by the speedometer. Likewise, the visual system can be said to have internalized some world knowledge if its behavior is smart.

In sum, to fulfill the criterion of internalization, a regularity has to be nontrivial and must have a chance to be ignored. Only these cases are subject to empirical testing. A second requirement, not explicitly imposed by Shepard, is best described in terms of Aristotle's classification of causes. If the regularity reflects knowledge of efficient causes in the environment (*causa efficiens*) it can be internalized. If instead the regularity reflects other knowledge, such as action goals (*causa finalis*), it cannot be internalized. Thus, to support the internalization hypothesis, we have to witness the use of rules that are finely resolved and that reflect world knowledge. Since there are many counter-examples at the high resolution end, and a good deal of arbitrariness at the low resolution end, is the internalization hypothesis valuable at all?

4. Externalization rather than internalization?

The answer to this question can still be positive if, apart from empirical evidence, we have another method to assess the fruitfulness of the internalization hypothesis. I posit that we do. A thought experiment that supposes an opposite principle may be able to generate important insights and help us decide whether we want to retain the internalization idea. I suggest externalization as the opposite principle. If this opposite principle leads to predictions that are clearly erroneous, we are likely to be on the right track with internalization. If, on the other hand, we apply the cumulative empirical evidence for internalization to the externalization hypothesis and it fares as well or better, then there is something wrong with internalization.

Principles guiding our perception in cases where the percept is under-specified may be a projection of our own body dynamics onto the perceived reality. In other words, *externalized* aspects of the motor system rather than internalized aspects of the physical world outside ourselves may provide defaults for our perceptions. This would require a logic opposite from that of internalization. The logic behind the idea of internalization is that some laws governing the universe have been generalized and incorporated into the visual system. An externalization logic, in contrast, would not focus on the receptive visual system but on the active motor system. Considering that the visual system has evolved to guide action and not to give us nice pictures of the world, this appears equally plausible. In a way, the visual system might have "internalized" features of its own motor system. Evidence for this route could be derived from the many instances of ideomotor-action, which demonstrate a very close link between the two systems (see e.g., Prinz 1987). The motor system, in turn, has of course evolved under the constraints that act in its terrestrial environment and therefore exhibits many features that fall under the realm of classical mechanics. The important difference lies in the fact that the motor system is action-oriented and generates its own forces. If default interpretations performed by the visual system are mediated by the action-oriented motor system, a very different set of laws might smarten the visual system. These laws are not abstracted versions of Newton's motion laws but, rather, abstracted versions of force-producing body mechanics.

Thus, let us consider how the visual system turns ill-posed questions into well-posed ones: it derives its default solutions from implicit knowledge of the motor system that it serves rather than from abstract universal laws that have observable consequences. This process might best be referred to as an externalization of body mechanics. An example that illustrates and empirically supports this idea is reported below. It concerns the understanding of ballistic trajectories of projectiles, which can be traced from Aristotle to our times.

4.1. Projectile trajectories

Not only our explicit intuitive knowledge about the dynamics of moving objects (Shanon 1976), but also our perceptual knowledge about these events is often erroneous – albeit to a lesser extent (Kaiser et al. 1992). For example, explicit knowledge about trajectories of falling objects is seriously flawed. Similar to the above-mentioned belief that a marble rolled through a C-shaped tube preserves its curvilinear impetus, many people believe that objects dropped from a moving carrier fall straight down, as if they lose their horizontal velocity component (McCloskey et al. 1983). Correspondingly, adult subjects do not favor the correct parabolic trajectory over other paths.

Hecht and Bertamini (2000) presented drawings of various possible and impossible trajectories of a baseball thrown over a large distance. Subjects favored a sine wave path over the parabolic trajectory, and generally all paths of some continuous curvature were judged to be fairly natural. The canonical trajectory shape was neither preferred nor singled out as special. This lack of perceptual understanding, if not evidence for kinematic geometry, might explain why beliefs about the shape of ballistic trajectories were rather warped from Aristotle's times through the middle ages. It was not until the early seventeenth century that Galileo suggested the correct parabolic shape (Wunderlich 1977) that ensues when neglecting air resistance.

In 1572, Paulus Puchner devised an interesting analysis of ballistic trajectories to instruct canoneers of the Saxonian artillery, as visible in Figure 5. Puchner was the weaponry expert at the court of the Saxon elector August. Puchner based his state-of-the-art prediction of cannonball trajectories and distances on the Aristotelian notion of a three-step flight path (Wunderlich 1977): a straight ascension phase followed by a curved arc phase, and a straight vertical drop. This three-step trajectory is not easily compatible with medieval impetus theory, because the circular phase of the flight path cannot be explained by air resistance diminishing the original impetus but only by gravity (or something else) continuously acting on the cannonball. The last straight-down phase was probably an empirical observation that cannonballs tended to drop from almost straight above.

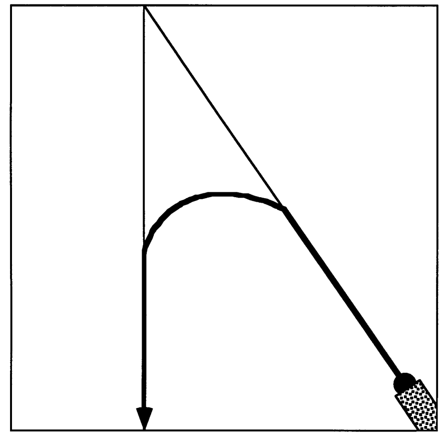

Figure 5. Ballistic trajectories as devised by Paulus Puchner in 1572. His state-of-the-art prediction of cannon ball shot distances was based on the Aristotelian notion of a three-step flight path. First a straight ascension path, second a circular arc, and finally a straight vertical drop. The first and last steps always had to complete a triangle at a given height when extended.

Notwithstanding these conceptual errors, observers have some visual knowledge about the correct parabolic trajectory and even better productive knowledge, as Krist and colleagues (Krist et al. 1993; 1996) have demonstrated. Their moving observers had to hit a stationary target on the ground by dropping an object. Given this facility, adult observers could easily have "internalized" the fact that cannon balls or rocks reach their maximal velocity when they exit the gun barrel or the pitcher's hand. The horizontal velocity component decreases as a function of drag and its change typically remains small in comparison to the deceleration of the vertical velocity component. The latter first decreases to 0, then the ball gains vertical acceleration on the downward part of its trajectory.

Surprisingly, many subjects believe that a ball will continue to accelerate after it has left the pitcher's hand. This belief is mirrored in perceptual judgments when impossible accelerating ball throws were presented in computer animation (Hecht & Bertamini 2000). Figure 6 depicts a trajectory that was judged to be rather natural. The cross indicates the point on the trajectory where observers believed the ball to possess maximal velocity. While these conceptual and perceptual data are grossly incompatible with any law of classical mechanics, including medieval impetus theory, they accurately describe the movement of the pitcher's arm. The latter does accelerate the ball and this accelerating movement might be projected onto the further movement of the projectile. These findings suggest that observers judge the throwing action as a whole and fail to parse the motor action from the mechanical event, or in other words,

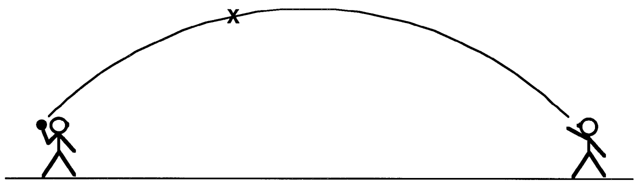

Figure 6. "Consider the path of a ball thrown from a pitcher (on the left) to a catcher (on the right). Mark the point where the ball has maximal speed." When asked this question, subjects' averaged answers correspond the position marked with a cross.

that the body mechanics of throwing have been externalized and projected onto the inanimate projectile.

A different interpretation may consist in the failure to conceive of the ball as an inanimate object. Obviously, as soon as the ball is no longer a projectile but has its own means of propulsion, an indefinite number of acceleration and deceleration patterns are compatible with the laws of physics. This interpretation is, however, rather unlikely given the evidence that even preschoolers are able to distinguish animate from inanimate objects (Massey & Gelman 1988). The dynamic motion context together with the mature age of the observers used in the above examples suggest that the ball was taken to be an inanimate object in these situations (see Gelman et al. 1995; Kaiser et al. 1992).

Even if the notion of externalized body dynamics appears to be rather speculative at this point, the misconceptions as well as the perceptual judgments reflect, in some sense, a continuation of a completed motor action (the acceleration of the arm) and an anticipation of a future motor action (deceleration of the arm and catch of the ball). This piece of evidence, at the least, demonstrates that there are potential competitors for the concept of internalization. If internalization is understood as a principle of abstraction that is prevalent in situations of underspecified perception, it must be legitimate to extend it to the realm of intuitive physics where perception and cognition overlap. In the above example, observers do not behave as if they have internalized regularities about projectile motion. An externalization account fits the data much better. Thus, our final attempt to provide support for internalization by comparing it to its opposite has failed. It even looks as if the notion of externalization has to be taken seriously in its own right.

5. Conclusion

I have tried to put the notion of internalization to a test while factoring in as many interpretations of the concept as possible. I have argued that internalization, interpreted in a narrow sense, is false. At the same time, broader interpretations run the risk of making the concept immune to any attempts to test it empirically. The middle ground is what deserves discussion. I have scrutinized this middle ground by recourse to examples from intuitive physics for two reasons. First, they are true to Shepard's spirit of looking for the relation of perception to the laws of physics. Second, they could logically have been internalized. I found mostly problems and counterexamples. Observers do not behave as if they make use of knowledge about gravitational acceleration, or the law of horizontal liquid surfaces. Neither do geometric geodesics prescribe our perceptual solutions in more than a few specialized cases. Internal knowledge about world regularities seems to be specific and task dependent, not universal.

5.1. Internalization versus habit

Shepard's concept can be understood as the phylogenetic complement to James's notion of habit. A habit, according to James (1890/1950), is a law that the organism has acquired during its lifetime and it thus has a clear ontogenetic character. Habits necessarily disqualify as internalized because they are acquired and can be changed. At the abstract level of analysis this appears to be acceptable. At the level of concrete examples, however, the distinction between habit and internalized rule is very hard to make. Take the case of a simple reaching movement. In a force field that is not typical for the gravitational field on earth, for instance when being spun on a centrifuge, observers do not take the unexpected forces into account. Their reaches are perturbed. However, after a few more reaches they have adapted to the unusual forces acting on their arms (Lackner & Dizio 1998). Does it make a difference in this case whether we say that the observer has formed a new habit or that she has quickly overcome the internalized regularity? Or has the normal gravitational force field (1 g) only been internalized when people fail to adapt to the new environment? Maybe resistance to adaptation can be used as measure for internalization. By spelling out the differences between habits and internalized regularities in such exemplar cases, the latter concept might be sharpened.

Two aspects of the visual system that are well captured by habit seem to render Shepard's approach cumbersome. First, the visual system is flexible. As in the water-level task, observers may change their behavior dramatically with experience. Experienced waitresses and bartenders produced larger deviations from horizontality than naive participants. Does that mean that internalized regularities can be modified on the spot? If this were the case, it would be almost impossible to differentiate between habits and internalization.

Second, Shepard's model excludes the natural environment as the major player in shaping the percept. The examples that I have discussed here attempted to focus on natural viewing situations. They failed to support his model. A system of bounded rationality (Simon 1969; 1990) such as the visual system may confine its solution space not by resorting exclusively to internal laws but rather by including environmental "satisficing" constraints that come to the fore depending on the environment and the situation in which the actor finds herself. This position could easily be extended to include the body dynamics of the actor as additional constraints.

5.2. Can the concept of internalization be salvaged?

Shepard's concept can deal in four ways with the failure to find evidence for regularities such as gravity and horizontality. First, all those potential invariants that did not pass empirical testing could be explicitly excluded from the theory. This would narrow the concept of internalization to a small set of applications.

Second, only certain regularities could be predetermined to qualify for internalization. This solution is also unacceptable. Thus far, Shepard has not provided any rules to decide when a regularity needs to have been internalized. As long as these are missing, the notion of internalized constraints does not have the status of a theory. It cannot be falsified. Shepard's statement could be formulated as: "Some regularities of the physical world have been internalized and act as constraints on perception and imagery." This is an existence statement which can only be disproved if we fail to find a single supporting case. Unfortunately, existence statements by themselves do not allow any predictions about other cases.

Third, another unacceptable salvage attempt would be to push the degree of abstraction of the concept even further. In some abstract sense, it cannot be wrong to claim that we have phylogenetically internalized some regularities of the physical world.

> When formalized at a sufficient level of abstraction, mental principles . . . might be found . . . perhaps attaining the kind of universality, invariance, and formal elegance . . . previously accorded only to the laws of physics and mathematics (Shepard 1994, pp. 2–3).

However, this venue of hyper-abstraction leads to immunity and removes internalization from the empirical discourse.

Finally, the only acceptable solution I see is to make the concept of internalization more powerful by adding specific hypotheses that rule out alternative explanations, such as ontogenetic learning of the circadian rhythm. Given the structure of Shepard's argument, such hypotheses should be derived from evolutionary theory. They might add some of the required resolution to the debate. It remains to be seen whether such salvage attempts are going to be worth the effort.

As an alternative to the exhaustion of salvage attempts, other competing concepts to internalization need to be taken seriously. I have assessed whether support for internalization could be derived from the failure of its opposite: the notion that we have not internalized world knowledge but externalized volitional and motor aspects of our own organisms. This opposite – externalization – fared quite well and merits further exploration.

ACKNOWLEDGMENTS

I wish to thank the Forschungsgruppe "Perception and Evolution" at the ZiF for raising the topic of internalization. I am indebted to William Epstein, Jessica-Gienow Hecht, Mary Kaiser, Michael Kubovy, Donald Hoffman, John Pittenger, Dennis Proffitt, Robert Schwartz, and four anonymous reviewers for valuable comments and suggestions.

NOTE

1. I realize that Chasles' theorem does not apply here. Kinematic geometry may not have a clear prediction for this case. However, a general default of curved movement would. And as we have seen, the degree of abstraction appropriate for internalization is highly debatable.