

suit of them has yielded impressive results. A question remains, however, one that falls outside the province of Shepard's discussion. The question concerns what exactly Shepard's equations purport to describe. Shepard speaks of "representations" and "appearances." This implies that what Shepard is after is a set of principles governing creatures' manipulation of representational states of mind.

Physicists employ equations to represent and explain the actions and powers of material bodies: bodies "obey" or "follow" laws these equations express. In just the same sense, intelligent creatures could be said to obey or follow laws of physics. Compare this to a case in which you obey a rule for stopping at stop signs by halting your car at a stop sign. Invoking a distinction made famous by Kant, we can say that you are guided by or *act on* the stop sign rule. In contrast, although your actions *accord with* laws of nature, you do not act on those laws. This is just to say that actions can accord with a law or principle without thereby being based on or guided by that law or principle. In acting on a principle, an agent's grasp or representation of the principle (in concert with other states of mind) controls the action.

What of **SHEPARD's** principles? Suppose Shepard has it right: creatures' assessments (explicit or implicit) of certain features of their environment conform to the principles he advances. Do these principles guide creatures' assessments of colors, or shapes, or motions? That is one possibility. Another, less ambitious, possibility is that creatures' actions merely satisfy the principles. If that were so, then the creatures need only possess a nature the physical composition of which supports mechanisms whose operation is describable via the principles.

Is this one of those philosophical distinctions without a difference? Certainly anything any creature does, if governed by any law, is governed by (and so accords with) basic physical law. Just as this need not be taken to imply that every science is reducible to (in the sense of being replaceable by) physics and chemistry, so it need not mean that explanations that appeal to principles on which agents are taken to act, are replaceable by explanations framed in terms of laws to which agents' actions merely conform. In invoking representations in explanations of creatures' actions we appeal to this very distinction.

Representing our surroundings differs from simply mirroring those surroundings. Representation is selective and partial; we represent the world in a particular way from a particular point of view. Evolution ensures that perceptual representations are constrained by the world. Our finite nature imposes additional constraints. This can be made to sound trite: the way the world looks, feels, sounds, and tastes to us depends on how the world is and how we are. But the formulation of principles that capture the workings of this mechanism is anything but trite.

Psychological explanation is susceptible to a peculiar sort of mis-direction. Features of the explanatory apparatus are easily mistaken for features of what is being explained. This occurs in everyday life when we anthropomorphize pets, ascribing to them states of mind they are unlikely to be in a position to harbor. Psychologists risk a similar confusion in formulating principles taken to govern mental representations. It is easy to mistake features of the formulation for features of the system. An example of a mistake of this kind might be the imputation of a mechanism for solving differential equations in the brain of an outfielder pursuing a fly ball. We describe the ball's trajectory using differential equations, and the outfielder's brain must incorporate mechanisms that arrive at comparable solutions. But it need not follow that the brain engages in computations of the kind we would use to describe the flight of the ball. Instead, the brain might avail itself of simpler heuristic mechanisms. One way to describe these mechanisms is to describe their inputs and specify a principle that takes these into appropriate outputs. But we cannot move directly from such a description to the conclusion that the brain operates on, and not merely in accord with, these principles.

Perhaps the nervous system is a "smart machine" or, better, a vast system of smart machines (Heil 1983; Runeson 1977). Smart machines are devices that execute computationally sophisticated tasks

in mechanically simple ways. A centrifugal governor on a steam engine is a smart machine, as is a polar planimeter (a simple device used to determine the area of irregular spaces, the area of an island on a map, for instance). Such devices act in accord with certain mathematical rules, but not on the basis of those rules. Knowing the rules would not tell you how the devices were constructed, how they actually operate. It is hard to avoid the impression that **SHEPARD's** principles are like this. In representing the world, we (or our visual systems) act in accord with these principles, but not on them. This is where talk of mental representations stands to be misleading. Mechanisms underlying the production and manipulation of our worldly representations could well operate in accord with certain principles without those principles mirroring the underlying mechanisms. Mechanisms operating in accord with the very same principles could well differ internally in important ways.

None of this affects the validity or significance of **SHEPARD's** results – results which, in any case, a philosopher is in no position to challenge. It does, however, affect the ways we might seek to understand and test those results in looking at the underlying hardware.

ACKNOWLEDGMENTS

The author is indebted to Davidson College for funding a research leave during 2000–2001, and to the Department of Philosophy, Monash University, for its hospitality, intellectual stimulation, and support.

What is the probability of the Bayesian model, given the data?

Evan Heit

Department of Psychology, University of Warwick, Coventry CV4 7AL, United Kingdom. E.Heit@warwick.ac.uk
<http://www.warwick.ac.uk/staff/E.Heit/>

Abstract: The great advantage of Tenenbaum and Griffiths's model is that it incorporates both specific and general prior knowledge into category learning. Two phenomena are presented as supporting the detailed assumptions of this model. However, one phenomenon, effects of diversity, does not seem to require these assumptions, and the other phenomenon, effects of sample size, is not representative of most reported results. [TENENBAUM & GRIFFITHS]

The Bayesian model proposed by **TENENBAUM & GRIFFITHS** has a number of strengths, such as extending Shepard's (1987b) account of generalization to multiple stimuli. This model is by no means the only model of categorization that extends **SHEPARD's** work (see, e.g., Nosofsky 1988b), nor is it the only Bayesian model of categorization to be applied to psychological data (see, e.g., Anderson 1991). Perhaps what is most important and novel about this modeling effort is the explicit emphasis on how people's prior beliefs are put together with observed category members to make classification judgments.

Since Murphy and Medin (1985), there have been many theoretical arguments and empirical demonstrations showing that categorization must be constrained by prior knowledge and cannot simply depend on generalization from observations (see Heit 1997b, for a review). However, model-based research in categorization has lagged behind on this important issue, with most categorization models not addressing influences of prior knowledge. In contrast, the Bayesian model of **TENENBAUM & GRIFFITHS** gives an elegant account of how two kinds of prior knowledge are incorporated into categorization.

First, category learning is set against the backdrop of a hypothesis space, which represents expectations about the possible content of the category. Category learning can be viewed as elimination of hypotheses that do not fit the data while strengthening the remaining hypotheses (cf., Horwich 1982). The Bayesian method for deriving posterior probabilities of hypotheses embodies the idea that not only does prior knowledge serve as a guide to what

the observed category members will be like, but also the observations themselves are crucial for selecting from among numerous prior hypotheses (Heit & Bott 2000).

Second, the modeling framework can apply general knowledge about how observations are sampled. This is knowledge not about the possible content of the category to be learned, but rather about the manner of learning itself. The crucial idea introduced by **TENENBAUM & GRIFFITHS** is “strong sampling,” an assumption that observations are drawn randomly from some fixed population. Strong sampling has important consequences, such as favoring specific hypotheses corresponding to smaller populations of positive examples – this is called the “size principle.” Within this modeling framework it could be possible to build in further distinctions about sampling, such as whether sampling is with or without replacement (Barsalou et al. 1998) or whether the observations have been presented in some purposeful order according to goals of a teacher (Avrahami et al. 1997).

In support of the Bayesian model including the size principle, **TENENBAUM & GRIFFITHS** focus on two phenomena, that more variable or diverse observations lead to broader generalizations, and that as the number of observations within a given range increases, generalization outside the range is reduced. These two phenomena are now considered in turn.

First, although the effect of diversity does appear to be robust, there have been salient exceptions reported in inductive reasoning tasks (reviewed by Heit 2000). Some cross-cultural work and developmental research has failed to find the diversity effect. Even with American college students, Osherson et al. (1990) reported an exception to the diversity effect: People draw stronger inferences given an observation that flies, for example, have some characteristic, compared to being given an observation that both flies and orangutans have this characteristic. It would be a challenge for any Bayesian account of induction, including Heit (1998), to address these exceptions, because Bayesian accounts do seem to predict robust diversity effects.

It is notable that Heit’s (1998) Bayesian model of inductive reasoning predicts diversity effects without any size principle or assumption of strong sampling. Indeed, use of information about variability of evidence is taken to be a hallmark of Bayesian models in general (Earman 1992). Likewise, models of categorization without any size principle, such as Nosofsky’s (1988b) exemplar model and Ashby and Gott’s (1988) parametric model, also predict broader generalization from more variable observations. Although it is clear from **TENENBAUM & GRIFFITHS**’s Figure 2 that it is possible to predict the diversity effect with strong sampling and the size principle, it seems that the diversity effect in itself is not strong evidence for these assumptions. Other models without these assumptions can also predict this result.

The second, fascinating result, is that with other things being equal, larger samples tend to promote less broad inferences (reported in Tenenbaum 1999). This result does seem to be distinctive evidence for the size principle, as illustrated by **TENENBAUM & GRIFFITHS**’s Figure 3. This result would not be predicted by categorization models without the size principle such as Nosofsky (1988b) and Ashby and Gott (1988).

However, this result differs from numerous results showing just the opposite, with larger numbers of observations leading to broader generalizations. Although it is hard to perfectly eliminate confounds between number of observations and their variability, it appears that Homa et al. (1981) did show greater generalization to categories with more members. Nosofsky (1988b) showed that when a category member is presented a large number of times, there is increased generalization of similar stimuli to the same category. Maddox and Bohil (1998) showed that people can track the base rates of categories, with a bias to put transfer stimuli in more categories with more members. None of these results are insurmountable evidence against **TENENBAUM & GRIFFITHS**’s Bayesian model; for example, Bayesian models can easily incorporate information about base rates. Yet, it does appear that the result presented by **TENENBAUM & GRIFFITHS**, that larger samples lead

to less broad generalization, is not characteristic of most results reported in this area. It would be important to establish the boundary conditions for this fascinating but isolated result.

In sum, the Bayesian model of generalization proposed by **TENENBAUM & GRIFFITHS** makes substantial contributions beyond existing accounts. However, the value of this model surely will be in its ability to address already documented phenomena in generalization, categorization, and inductive inference, including the exceptions to the diversity and sample size effects predicted by the model. In the target article, the model is applied to tasks where only positive cases of a single category are presented. Although it is valuable to focus on this important learning situation, it is notable that many more psychological experiments have addressed learning to distinguish one category from another, or learning from positive and negative examples. To address this large body of existing research, the Bayesian model itself would require some further generalization.

Adaptation as genetic internalization

Adolf Heschl

Konrad Lorenz Institute for Evolution and Cognition Research, University of Vienna, A-3422 Altenberg, Austria. Adolf.Heschl@klu.univie.ac.at
<http://www.univie.ac.at/evolution/kli>

Abstract: In the course of evolution organisms change both their morphology and their physiology in response to ever-changing environmental selection pressures. This process of adaptation leads to an “internalization,” in the sense that external regularities are in some way “imitated” by the living system. Countless examples illustrate the usefulness of this metaphor. However, if we concentrate too much on Shepard’s “universal regularities in the world,” we run the risk of overlooking the many more fascinating evolutionary details which alone have made, and still make possible the evolution of diversity on earth.

[SHEPARD]

I will first attempt to theoretically underpin the concept of “internalisation” as it has been used and further developed in an impressive way by Roger **SHEPARD**. Let us begin at the lowest imaginable level of evolution: the genetic modification of organisms through random variation and natural selection. As has been empirically shown (Luria & Delbrück 1943), random genetic variation forms the molecular basis for subsequent evolutionary processes. These mutations can have three effects on the biological fitness of their carriers: they can be neutral, negative or – rarely – positive. In the first case, the traits of the organisms concerned will vary in a completely random way and no structuring effect of the environment will be recognizable. In the second case, many or even all carriers of the mutation will ultimately disappear from the evolutionary scene. This will have a clear structuring effect on the whole population of a given species, in the sense that only those individuals lacking this mutation will survive and reproduce. In such a case, we could speak of negative selection or, more simply, an extermination effect of specific adverse external influences.

The third case, in which a new mutation provides an advantage for the organisms concerned, leads to something we could indeed call an “internalization” of external regularities. To give an example at the molecular level: an enzyme (lactase) is produced which allows humans to better digest the form of sugar found in milk (lactose); as is to be expected, the distribution of the gene coding for this enzyme within the population reflects the structure of a concrete external regularity: the geographic distribution of intensive dairy-farming (Jones 1992). What is valid at the population level must also be valid at the molecular level: the chemical structure of lactase in turn reflects certain specific structural aspects of the disaccharide lactose. Hence, in an evolutionary perspective, it is perfectly legitimate to equate the process of biological adaptation with a kind of internalization process of external selection pressures, because *every* adaptive change must necessarily be accompanied by a corresponding form of internal restructuring.