tion.) SHEPARD rightly acknowledges that we have to assume a rich internal structure of the perceptual system in order to account for the relevant facts. He thus draws our attention to a core problem of perception theory, viz., to understand the structural form of internal representations. To this end, SHEPARD extends the approach of ecological physics to further kinds of abstract mathematical descriptions of external regularities, which he then uses as heuristics for exploring the structure of internal representations. His grand perspective on the *Evolution of a mesh between principles of the mind and regularities of the world* (1987a) doubtless includes a good deal of truth, notwithstanding the problems that his notions of "regularity" and "internalisation" are faced with when one attempts to understand them beyond their meaning in ordinary discourse. SHEPARD's more extensive (non-Darwinian) claim that there is an "evolutionary trend toward increasing internalization" (1987a, p. 258) and that by internalising more and more physico-geometrical regularities the fitness of a species is increased, is not easy to assess and would hardly be maintained in other areas of biology. Fortunately, issues of evolutionary internalisation do not bear any immediate relevance with respect to perceptual theory, because here, as elsewhere in biology, a satisfactory *ahistorical* account for a functional structure does not ipso facto suffer from some kind of explanatory deficit (cf. Fodor 2000). It seems to me that the role that the concept of internalisation plays in SHEPARD's account resembles the role that mechanisms of association play in standard empiricist approaches to the mind, viz., it acts as a kind of general multi-purpose acquisition device for building up mental structure.

What appears to me to be more problematic than the metatheoretical discourse about internalisation is **SHEPARD**'s extreme physicalistic stance. In SHEPARD's view the structure of internal representations is determined predominantly by regularities of the external world, whereas no essential explanatory importance is attached to those aspects of the internal conceptual structure of perception that do not mirror external regularities, or to internal constraints of the cognitive architecture. Shepard (1984, p. 431; 1987a, p. 269) seems to think that constraints on the principles of the mind that do not have an external origin are merely arbitrary. Naturally, they must appear arbitrary if one slices the nature of perception according to external physical regularities, thus succumbing to the *physicalistic trap* in perception theory (cf. Mausfeld, in press).

Evidently, there is sufficient overlap between regularities of the world and the structure of internal representations. However, from this global property, which pertains to the entire organism, it does not follow that the representational structure of specific subsystems is predominantly determined by specific features of the environment. With respect to internal codes, equivalent classes of sensory inputs are held together by the conceptual structure of our perceptual system, rather than by the structure of the physical environment itself. The given conceptual structure that is part of our biological endowment is based on concepts that are not expressible as "natural kinds" or abstract regularities of the external physical world. This is evident for internal perceptual concepts such as "edible things and nutrients" or "emotional states of others." In other cases, such as the *internal* concept "surface colour," it may be less obvious that it defies definition in terms of a corresponding physical concept (even in the sense of the latter providing necessary and sufficient conditions for the former). Rather, it has its own peculiar and yet-to-be identified relation to the sensory input and depends intrinsically, in an idiosyncratic way that cannot be derived from considerations of external regularities, on other *internal* codes, say, for perceived depth or figural organisation. (All the same one might be able to concoct some Panglossian post hoc story in terms of external regularities for each specific case, but nothing about an external origin would be implied by this.) The structure of internal representations, as Gestalt psychology and ethology have already provided ample evidence for, is shaped not only by regularities of the external world. Rather, internal representations have to fit into the entire conceptual structure of the perceptual system, including its two fundamental interfaces, viz., the interface with the motor system and that with the higher cognitive system, where meanings are assigned in terms of "external world" properties.

SHEPARD has reinvigorated psychological inquiries into the structural form of mental representations. Such inquiries inevitably lead back to the core problem of perception theory, viz., to understand the internal conceptual structure of perception. This problem, however, cannot be solved or dodged by exclusively referring to physico-geometrical or statistical regularities of the external world and by assuming that the rich structure is imprinted on the mind of the perceiver almost entirely from without.

While **SHEPARD** seems to accept internal structure only to the extent that an external origin dignifies it with a stamp of approval, as it were, **KUBOVY & EPSTEIN** relapse altogether into a wariness about postulating specific internal structures. They refer to a distinction, widespread in empiristic approaches to the study of the mind, between what they call a "measurement model of perception" and assumptions of "invisible internal principles." Because they do not want to lodge the principles that are part of a successful explanatory account in the mind of the percipient, they propose what they call a "more modest interpretation," according to which we can, instead of talking about internal principles, only say that the visual system proceeds *as if* it obeys internal principles. Thus, they implicitly make the distinction between evidence for an explanatorily successful theory and evidence for the "psychological reality" of the principles to which this theory refers. Even if SHEPARD's investigations on motion perception provided, at the level of description on which he is working, a successful explanatory account – both in range and depth – of an important class of facts, it would still lack, in KUBOVY & EPSTEIN's view, "psychological reality." This is a highly questionable and unjustified distinction, which would hardly be of interest elsewhere in the natural sciences. A similar request for a "more modest interpretation" in physiology with respect to, say, the idea that "pattern cells in area MT employ the assumption of smoothness in their computations of motion" (Hildreth & Koch 1987, p. 508) would justly be regarded as being without any theoretical interest. In perception theory, as in other fields of the natural sciences, we proceed by attributing to the system under scrutiny whatever serves our explanatory needs. Ascribing inner structure to the perceptual system is not some mysterious ontological commitment, but a case of an inference to the best explanation (subject to further inquiry and open to change). There are (aside from metaphysical issues) no ontological questions involved beyond what is stated by the best current explanatory account. The distinction that KUBOVY & EPSTEIN make is an instance of what Chomsky (2000) has identified as an odd dualism of explanatory principles between psychology and the rest of the natural sciences. Such a dualism, which is expressed in KUBOVY & EPSTEIN's emphasis on "measurement theories of perception," will impede asking, as SHEPARD does, fruitful questions about the "invisible internal principles," – a natural concern, it seems, for inquiries into the nature of the perceptual system.

Probabilistic functionalism: A unifying paradigm for the cognitive sciences

Javier R. Movellan and Jonathan D. Nelson

Department of Cognitive Science, University of California San Diego, La Jolla, CA 92093. **movellan@mplab.ucsd.edu**
http://www.mplab.ucsd.edu/ jnelson@cogsci.ucsd.edu
http://www.cocsci.uscs.edu/~jnelson/

Abstract: The probabilistic analysis of functional questions is maturing into a rigorous and coherent research paradigm that may unify the cognitive sciences, from the study of single neurons in the brain to the study of high level cognitive processes and distributed cognition. Endless debates about undecidable structural issues (modularity vs. interactivity, serial vs. parallel processing, iconic vs. propositional representations, symbolic vs.

connectionist models) may be put aside in favor of a rigorous understanding of the problems solved by organisms in their natural environments. [SHEPARD; TENENBAUM & GRIFFITHS]

TENENBAUM & GRIFFITHS' (T&G's) paper on SHEPARD's law of generalization is a beautiful example of the most exciting and revolutionary paradigm to hit the cognitive sciences since connectionism. We call this paradigm "probabilistic functionalism" for its focus on functional rather than structural questions and for its reliance on the machinery of probability and information theory. Probabilistic functionalism traces back to Brunswik (1952) and finds modern articulators in Marr (1982), Anderson (1990), and Oaksford and Chater (1998).

Suppose an organism was rewarded for pecking in response to a red key. Would the organism generalize the pecking behavior in response to a purple key? Shepard (1987b) observed that in a very wide variety of experiments, the degree of generalization to new stimuli is an exponential function of the perceived similarity between the old and new stimuli (see Fig.1, top). Under the dominant structural paradigm, one would typically approach this result by formulating mechanistic models of information processing, for example, connectionist networks, that exhibit the law. Yet even under the unrealistic assumption that one can uniquely specify the mechanisms of the mind, the structural approach ultimately fails to answer a critical question: why does the mind use such mechanisms?

Instead of the dominant structural approach, Shepard (1987b) framed generalization as the reflection of a Bayesian inference problem: specifying the category C of stimuli that lead to a given consequence. In our example, C would be the set of colors that lead to the reward. SHEPARD assumed that the degree of generalization to a new stimulus y is proportional to $P(y \in C | X = x)$, the probability that y belongs to the category C, given the example x. He then found sufficient conditions under which this function is approximately exponential. The ubiquity of the exponential law is ultimately explained by the fact that these conditions are reasonable for a wide variety of realistic problems. T&G extend Shepard's analysis to cases in which some of Shepard's conditions are not met: they allow multiple examples and nonconvex consequential regions. In both cases, violations of the exponential law are possible (see Fig. 2, 3, and 5 of T&G's paper).

SHEPARD's argument makes no distinction between conscious and automatic inferences. In addition, it is gloriously silent about representational and processing issues. In the functional approach, probabilities are just tools used by scientists to understand conditions under which observed behaviors are reasonable. These probabilities do not need to be explicitly represented by the organism under study. Consider, for example, the classic signal detection problem of discriminating a known signal in the presence of white noise. One can implement an optimal classifier for this problem without computing any probabilities at all. All the system needs to do is to estimate the correlation between the observed data and the known signal and to make decisions based on whether such correlation is larger than a threshold. Still, a Bayesian functional analysis in terms of subjective probabilities will be useful to

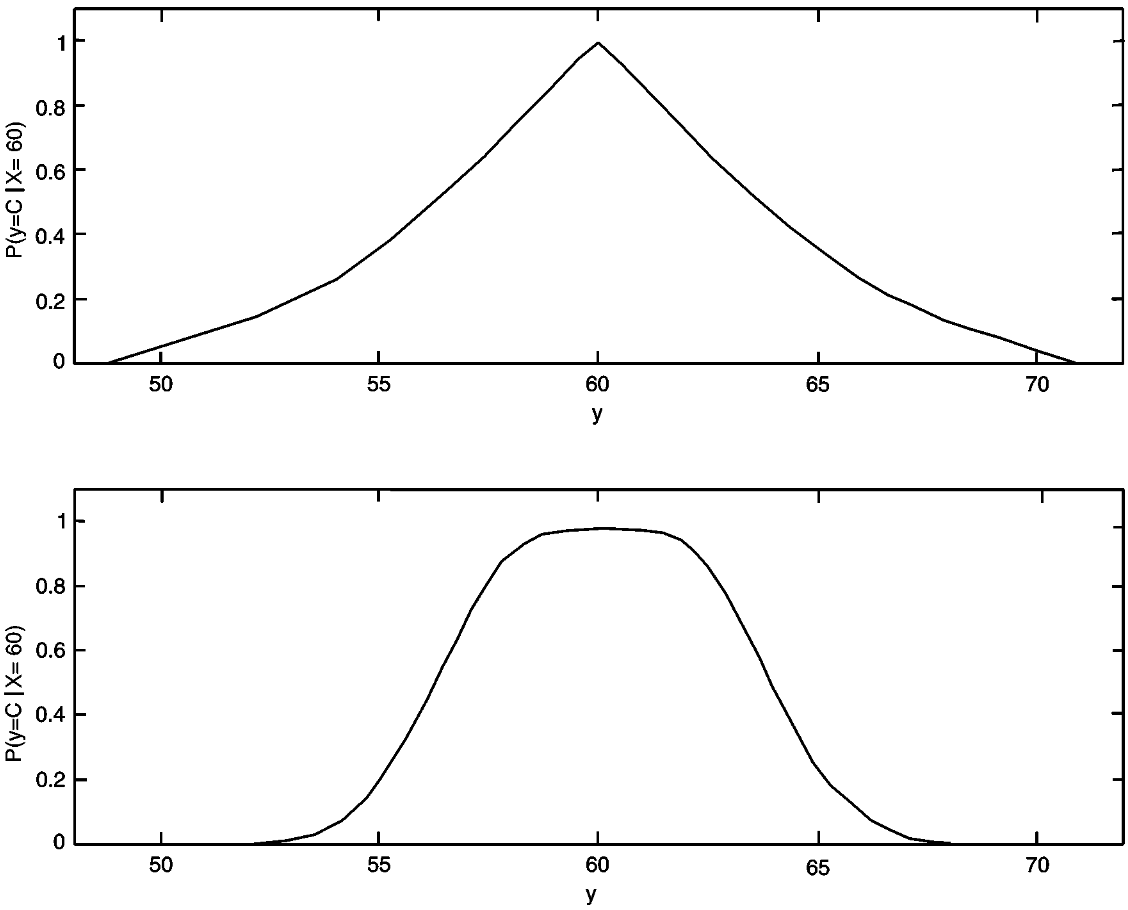

Figure 1 (Movellan & Nelson). Effects of non-uniform sampling on the generalization law. This figure shows the probability that a value y belongs to the interval C, given the fact that the value 60 was sampled from that category. Top: The distribution of examples from C is uniform, resulting in a concave upwards generalization law. Bottom: The distribution of examples from C is Normal centered at C and with standard deviation equal to 1/8 of the length of C. This results in a violation of the exponential law. In both cases the prior distribution for C was the same: uniform for location and positive truncated Gaussian for scale.

understand the conditions under which the system is optimal.

The fact that we do not need to worry about how probabilities are represented makes the functional approach easily portable to a very wide variety of problems: from the study of single neurons, to the study of perception, conscious decision making, and the study of distributed cognition. It is thus not surprising that the emerging success of probabilistic functionalism reaches across a wide range of disciplines in the cognitive sciences: probability theory has become the language of choice to understand computation in neural networks (Bishop 1995). Bell and Sejnowski (1997) showed that the receptive fields of simple cells in primary visual cortex are optimal for transmission of natural images. Lewicki (2000) used similar techniques to show the optimality of cells in the auditory nerve. Knill and Richards (1996) illustrate the power of Bayesian techniques to understand perception. The "rational" movement in cognitive psychology is a perfect illustration of how functionalism can help us understand high level cognition (Oaksford & Chater 1998).

Extensions of the analysis of generalization.

Shepard (1987b) provided a solid foundation for the functional analysis of generalization and T&G extended the analysis in important ways. However, there are still some outstanding issues that need to be addressed. In this section we focus on such issues.

Response rule. Consider the case in which an organism is rewarded for pecking a red key $\boldsymbol{x}$. The current analysis assumes that the rate of response to a novel key $\boldsymbol{y}$ is proportional to $\boldsymbol{P}(\boldsymbol{y} \in \boldsymbol{C} \,|\, \boldsymbol{X} = \boldsymbol{x})$, the probability that $\boldsymbol{y}$ belongs to the consequential region $\boldsymbol{C}$, given the example $\boldsymbol{x}$. While this function exhibits the desired exponential law, it is unclear why one may in general want to respond with a rate proportional to such probability. For a functional explanation to be complete, this point needs to be addressed.

Typicality and sampling. While T&G significantly generalize SHEPARD's analysis, they still constrain their work to the following conditions: binary membership functions and uniform distribution of examples within categories. These assumptions may be still too restrictive. It is well known (Rosh 1978) that humans do not treat all elements of a concept equally (e.g., robins are better members of the category birds than penguins are) and, thus, graded membership functions may be needed to model human inference. Moreover, in many situations, examples do not distribute uniformly within categories and this may result in significant changes of the generalization function. For instance, take the case used by T&G of a doctor trying to determine the healthy levels of a hormone. A healthy patient has been examined and found to have a hormone level of 60. What is the probability that another hormone level, for example 75, is also healthy? In this case the consequential region $\boldsymbol{C}$ is an interval representing the set of healthy hormone levels. If we assume that hormone levels have a uniform distribution within that interval then the exponential law of generalization follows (see Fig. 1, top). However, if we let hormone levels to be normally distributed within the interval, that is, more probable about the center of the category than at the extremes, then the exponential law can be violated (Fig. 1, bottom).

The size principle. According to the size principle proposed by T&G, smaller categories tend to receive higher probability than larger categories. Under the assumption that examples are uniformly sampled from categories, this is just a consequence of one of the axioms of probability. However, if examples distribute in a non-uniform manner within the category, for example, if robins are more likely to be sampled as members of the category "birds" than penguins are, the size principle would need to be reformulated, perhaps in terms of the entropy of the sampling distribution rather than the size of the category.

Statistical analysis of the environment. SHEPARD and T&G frame their analyses in terms of subjective probabilities. Thus, it is entirely possible for the generalization law to be subjectively optimal and objectively inadequate. We believe a crucial part of probabilistic functionalism is to analyze the statistics of actual environments and to test whether the assumptions made by functional models are reasonable for the environments at hand. See Movellan and McClelland (2001) for an example of how this analysis may proceed in practice. This issue needs to be addressed in the context of the exponential law of generalization.

Making predictions. Besides offering useful descriptive insights, probabilistic functionalism can also be predictive. For example, Movellan and McClelland (2001) analyzed a psychophysical regularity: the Morton-Massaro law of information integration. This law is observed in experiments in which subjects integrate two or more sources of information (e.g., the speech signal from a person talking and the visual information from the talker's lip movements). According to this law, ratios of response probabilities factorize into components selectively influenced by only one source (e.g., one component is affected by the acoustic source and another one by the visual source). Previous debates about this law centered on a structural issue: Is this law incompatible with interactive models of perception? Movellan and McClelland found that both feed-forward and interactive mechanisms can perfectly fit the law and thus this structural issue is undecidable. In contrast, a functional Bayesian analysis of the law helped find a novel task for which the Morton-Massaro law should be violated. Experiments confirmed this prediction. Similar predictive analyses would be helpful in the context of the generalization law.

Beyond an occult kinematics of the mind

Keith K. Niall

Defence and Civil Institute of Environmental Medicine, Defence Research and Development Canada, Toronto, Ontario M3M 3B9, Canada.
Keith.Niall@dciem.dnd.ca

Abstract: The evidence for a kinematics of the mind is confounded by uncontrolled properties of pictures. Effects of illumination and of picture-plane geometry may underlie some evidence given for a process of mental rotation. Pictured rotation is confounded by picture similarity, gauged by gray-level correlations. An example is given involving the depicted rotation of Shepard-Metzler solids in depth.
[HECHT; KUBOVY & EPSTEIN; SHEPARD; TODOROVIČ]

Brave explorers have often mistaken the nature of the greatness they discover. Columbus thought he had sailed to the Indies; Frege thought he had reduced arithmetic to logic; SHEPARD thought he had found the kinematics of mind. SHEPARD proposes that the "abstract constraints of geometry" (i.e., the three-dimensional geometry of our terrestrial environment) are separable from the "concrete constraints of physics." He proposes that the former – essentially kinematic – constraints match what is represented in vision and visual imagination, better than other constraints which are characteristic of physical dynamics. That is, he draws a dichotomy between kinematics and dynamics, and bases his theory of representation on kinematics. The dilemma may be premature, for there is more to consider in a theory of vision and visual imagination. Conditions of illumination, and the perspective geometry of pictures should also be subsumed; for, though illumination and perspective geometry may not seem central to the psychology of representation, they are central to the study of vision.

Sometimes explanation is simpler or closer to hand than one may imagine. The "mental rotation" effect may not concern kinematics or rotation in three dimensions at all; I argue it depends mainly or wholly on the perspective geometry of pictures. The bulk of this commentary is devoted to presenting a small illustration of this point for some depicted rotations in depth.

Consider the Shepard-Metzler solid that is depicted in Figure 1. A single solid composed of many cubes has been rendered in six perspective views: that is, in six perspective pictures. These perspective pictures represent rotations in depth: rotations of -120%, -60%, 0%, 60%, 120%, and 180% about a vertical axis