Philosophy and Phenomenological Research
Vol. LXXVI No. 2, March 2008

Seeing, Doing, and Knowing: A Précis

MOHAN MATTHEN
University of Toronto

I. The Sensory Classification Perspective

The theoretical fulcrum of *Seeing, Doing, and Knowing*[1] (*SDK*) is the *Sensory Classification Thesis*:

> Sensory systems assign distal objects to classes.

For example, colour vision assigns objects to colour-classes; form vision assigns them to shape-classes; speech perception assigns sounds to phoneme-classes; and so on. A *sense-feature* is a property that something has by virtue of belonging to such a class.

Sensory classification feeds into automated *epistemic operations*: for example, classical and operant conditioning (updated versions of Hume's *associations of ideas*) and habituation (diminished attention to repeatedly presented stimuli). These operations allow an organism to construct and update its records concerning the state of the world.

When sensory systems assign two stimuli to the same class, these stimuli are treated as the same with regard to automated epistemic operations; when they are assigned to different classes, they are treated differently. When two distal stimuli are co-classified, they elicit the same unconditioned responses. Similarly, when a stimulus is classified the same as others presented *earlier*, the perceiving subject tends (by the process known as "habituation") to pay it less attention. In this way, sensory states are simultaneously the *outputs* of sensory systems as well as *inputs* to conditioning and other automated operations.

Sensory consciousness serves as an indication of how the senses have classified something. A thing looks blue because once the sensory

[1] Mohan Matthen *Seeing, Doing, and Knowing: A Philosophical Theory of Sense-Perception* Oxford: Clarendon Press, 2005.

system has assigned it to that colour-class, it signals that it has done so by tagging it with a blue "look." This is the *Sensory Signalling Thesis*. The Sensory Classification *Perspective* consists of this thesis together with the Sensory Classification Thesis. (See *SDK* 30–35 for an elaboration of the idea that classification precedes appearance, and chapter 10 for the notion of sensory signalling.)

Sensory signalling is required only for some of the uses to which we put sensation. Some action-consequences of sensory classification are "coerced" (235–237)—conditioning, for example—and beyond the control of the perceiving organism. If I receive a shock after hearing a tone, I cannot help but flinch when I next hear a co-classified tone, even if I do not believe that I will be shocked. These coerced consequences do not require a conscious sensory appearance. Other action-consequences, however, are reflective and voluntary—when a thing *looks* blue, I may conclude that it *is* blue; or hang it on a yellow wall; or conclude from prior experience that my daughter will refuse to wear it. The blue look facilitates and is an essential intermediary for such *non*-coerced uses of sensory classification (237–239). In short, sensory consciousness is a medium for *entertaining* sensory data for reflective processing.

This all sounds very simple—and I hope it is—but it reverses the order of priority that has been standard in philosophy. Until quite recently, philosophers defined *being blue* in terms of *looking blue*. Things were blue if they had a propensity to look blue in normal circumstances. Or they *were* blue in situation *S* if they *looked* blue in *S*. The Sensory Classification Perspective goes the other way: a sensory system makes a thing look blue only after it has determined that it *is* blue. Sense-features such as *blue* or *round* must therefore be defined independently of the looks they present.

Recently, a new trend has tentatively emerged in philosophy that defines sense features independently of experience. Physicalism, which defines sense-features in terms of physical quantities (such as reflectance or wave-length) is one, but not the only, example of this trend. Physicalism leads to a puzzle. How can we have immediate knowledge of physical features? What kind of knowledge of sense-independent features do the senses provide? The Sensory Signalling Perspective is another example of this trend, but it is inoculated against the puzzle: the immediate epistemic response to sensory classification constitutes instinctive knowledge of sense-features from the subject's point of view. (In her comments, Frances Egan has an excellent summary of how I approach this problem in *SDK* chapter 11—see her discussion of the Fundamental Principle of Colour Attribution.)

II. Sense Features and Sense Orderings

Sensory classes are hierarchically ordered—each term in the series *ice blue*, *blue*, *cool-coloured* is contained in the next. That is, anything that is seen as ice blue is also seen as blue, and anything blue as cool in colour. The traditional view is that perceivers abstract such features from a concrete, non-articulated sensory image. Each term abstracts from the preceding, and so each presupposes the earlier terms in such a series. The Sensory Classification Perspective takes a different view: features such as the above are part of what visual sensation delivers.

The view expressed in *SDK* (70–77) is that many hierarchically distinct classifications emerge from *distinct* processes of sensory classification. The visual system delivers several distinct layers of colour classification—very broad classifications such as *light/dark* and *cool/warm*, narrow classifications that correspond to basic colour terms (*red*, *blue*, *yellow*, etc.), even narrower classifications about shades of colour, and so on. Many of these classifications emerge from separate sensory processes; they are not all derived from a single sensory image. For example, the broadest classifications depend on only a subset of colour-sensitive cells, and are available to colour-blind observers. But information about shades of colour involves fine grained opponent and contrast-sensitive processing. These are parallel processes; they deliver sense-features that are distinct from one another in somewhat the same way as *blue* and *round* (though of course there is more overlap between distinct colour processes than there is between colour and shape). A colour-blind person may sense the cool colour of a glacier without seeing its precise shade of ice-blue; the *wide* classification is delivered independently of the narrower one, and cannot therefore be held merely to be abstracted from the latter. The separation of seemingly overlapping processes is perhaps most intuitively evident in speech perception, where phonemes, pitch, timbre, and melodic contour are separately processed and then bound together. Someone could be deficient in the perception of any one of these without losing the others.

It is tempting to say that at the ultimate fineness of grain, sensory classes correspond to *determinate sense features*—in colour vision, for example, these would be *shades;* in shape perception, they would be determinate shapes; more generally, they would be sensory classes such that no two members can be discriminated from one another (with respect to colour, or shape, or whatever the determinable is). However, it is quite well known that this conception of determinate features is logically problematic because indiscriminability is not transitive—x could be indiscriminable from y and y from z, but not x from z. Yet sense-features should correspond to equivalence classes; thus, they cannot properly be defined in terms of visual indiscriminability.

This brings us to a more complex version of the Sensory Classification Thesis. This is the *Sensory Ordering Thesis*, the idea that sensory systems *order* distal stimuli in relations of similarity, rather than assigning them to discrete classes. The basic notion, according to this more complex point of view, is not *x is F* or *x and y are the same with regard to determinable D*, but rather *x is more similar to y than to z with regard to determinable D*. Operationally, such similarity relations are constructed by collating subjects' capacities to distinguish stimuli, and they are graphically represented in terms of "similarity spaces." Many intricate questions arise as to the psychological, neurological, and ontological reality, and the logic, of such representations. What is similarity? What is the significance of the dimensionality of similarity space? What is in-betweenness? Such questions are explored in Part II of *SDK*.

III. Sensory Classification: Some Corrolaries

The Sensory Classification Perspective contradicts the view that a sensation is simply the end-point of a physical process of information-transmission from an environmental object. The latter view suggests that physical similarity is carried from an object to a sense-organ by a natural process of transmission, thus causing similarity in sensations. Similarity of sensations thus indicates a similarity of the process of transmission. The Sensory Classification Perspective supposes instead that the similarity of sensations reflects the system's activity, which is directed toward detecting similarities in the world outside despite *dissimilarities* in the intervening causal chain.

Right from receptoral cells onwards, each stage of a sensory data-stream is devoted to registering relevant similarities and discarding irrelevant accompaniments. "Edge-detectors," for example, are devoted to finding lines in a retinal image across which there is an abrupt change of the level of illumination. They are indifferent to whether the brightly lit side is on the left or on the right of such a line, and what the colour of light happens to be; all such irrelevant data is simply thrown away. Notice, here, that a retinal image in which bright pixels are to the left of some line, and dark ones to the right is as dissimilar as can be, pixel by pixel, from one in which this polarity is reversed (48–49). Yet the edge-detector would *co*-classify these images; for they both incorporate an edge at the same place. Edge-detectors search for a feature that is buried in the flow of information. They *extract* information rather than merely receiving or recording it. Similarly, in audition, *phoneme-detectors* search for similarities in speech-production events, where these may be hidden in dissimilar acoustic patterns

(213–222). From the receptoral level on up, sensory systems consist of *feature-detectors* of this sort—centre-surround detectors, motion-detectors, Boolean calculators, and so on. Sense experience is a marker of features extracted by many layers of such detectors.

Understood in this way, the output of a sensory system signifies something in the form of a Fregean object-concept complex—i.e., a property attributed to a distal object. Sense features are functions from objects to truth-values; the task of a sensory system is to compute such functions, given receptoral activation levels as input. So it seems that the Sensory Classification Perspective contradicts a widely—indeed, almost universally—held thesis in contemporary philosophy of perception, namely that sensory content is *non-conceptual*. Chapter 3 of *SDK* responds to some main arguments that have been given in support of the non-conceptual content thesis.

IV. Coevolution and Content

As stated so far, the Sensory Classification Perspective is compatible with the view that sensory systems group things together in an arbitrary fashion, and that their classificatory scheme are utterly uninformative about the state and condition of the world outside. Our senses group green things together—but it could be for all that has been said so far that green things have nothing in common other than that the visual system groups them together. Evidently, though, biological sensory systems are products of evolution, and as such they must have been of some use to their owners. Does *this* suggest that green things do indeed share something? Perhaps: but this argument doesn't tell us much about what they share. Some have suggested that sense features such as colours or shapes are independent features of the world on which sensory systems converge. Others have disputed such an assumption.

Let's reformulate the problem. Suppose that two things look the same in some way—in colour, or shape, or size for instance. According to the Sensory Classification Perspective, these things are treated in the same way with regard to automatic epistemic operations. Consider an operation such as conditioning. Very crudely speaking, conditioning generates expectations concerning unobserved correlations. Suppose that an object x is placed in sensory class C. Suppose further that many members of C have been found to be F (for some F distinct from C), and at most a few to be non-F. Conditioning ensures that if some hitherto unobserved thing is found to be C, then it is expected to be F. Sensory classification is the precondition for these spreading expectations; it is because x is classified as C that the F-expectation spreads to it. So at the very least, sensory classification must evolve in such a way as to ground reasonable expectations.

This way of formulating what the senses do allows us to restate the assumption about evolutionary advantage. Suppose that an animal has encountered substances with smell *C* and found them to be sick-making. Conditioning leads it to expect that other substances with smell *C* will also be sick-making; and thus it avoids them. It is assumed that it will be advantageous in evolutionary terms that its sense of smell should classify things appropriately for such an inference. In other words, given that this is how it uses smell categories, it will be advantageous for it to co-classify things in such a way that things that fall into the same category of smell will belong to the same sick-making category. Thus, it is not that sensory systems *must* converge upon some pre-existing scheme of reality (though they may do this); rather sensory systems and epistemic operations *co-evolve* so that they are useful together (222–229).

With the above schema in mind, we might identify three kinds of property things share when they have smell *C* (assuming that the sensory system has performed properly). *First*, they share some real-world property characterizable independently of sensory systems—presumably, a chemical property which correlates with the property of being sick-making. *Second*, they have the property that they *ought* to be treated alike by various epistemic operations—for example, that the sick-making expectation should spread to them in virtue of their smell. *Third*, they have the property that they ought to be avoided because they might be sick-making.

Of these, the last seems to be a property that is in the realm of the *effector* mechanism, rather than that of the sensory system—it is a property that relates to how sensory output is used, not to the output itself. The second is proprietary to the operation of the sensory system itself, and expresses the internal significance of the system's output. I call this *primary sensory content* (233). The first, the chemical property, is arguably what the system detects—though it does not, as it were, know it under this system-independent description. I call this *secondary sensory content*. The relationship between primary and secondary content is analogous to that between sense and denotation (232–234).

V. Pluralistic Realism

Both primary and secondary sensory content are objective. It either is or is not the case that things that look the same in a certain way *ought* to be treated in the same way epistemically—where the 'ought' is understood as 'contributes to evolutionary advantage', or possibly 'contributes to the advantage of organisms like this'. Further, it is possible that in any given case, the classification is successful—the system may, for instance, classify something as having feature *F* and therefore dictate that it should epistemically be treated in the same way as other *F* things. However, it could be wrong—the thing might not really be the same as other *F* things. Thus, there are objective standards concerning the correctness of sensory classification, and we are entitled to a form of realism about sensory qualities.

This said, we need to recognize that since different kinds of organism use sensory information in different ways, they might sense very different features of the world. An extended argument as well as evidence to this effect is offered in Part III of *SDK*. The colours that birds sense should not be identified with those that primates sense; the phonemes that humans hear do not match up with what dogs or birds hear when they are exposed to the same sounds (213–222). During the course of evolution, each kind of organism arrives at a set of features that is useful to its style of life, given the kinds of sensory receptors available to it.

This is the thesis of *Pluralistic Realism*: different kinds of organism represent different features of the world, but there are objective standards of correctness with respect to the features each represents.

VI. Visual Reference

The first four parts of *SDK* are about sense-*features*, i.e., about the properties that sensory systems attribute to objects. Part V is concerned with how we make *visual* contact with external objects.

Chapter 12 is concerned with the kinds of objects that are involved in visual consciousness. It is often held that visual consciousness is of a *field*: a connected two- or three-dimensional array of *places* in which visual features reside. I contest this view. Visual awareness of *motion* implies that feature-instances can move from place to place while remaining in the same subject (272–282). This implies that visual features appear to be attributed to things that can move from place to place—material objects, in other words. Thus, visual consciousness is of objects not places; I argue that we are not visually aware of unoccupied places. This implies that visual awareness is *not* of a connected array—an anti-Kantian conclusion. Chapter 12 also discusses the objects of audition and olfaction—these are not material objects, but sounds and smells.

Chapter 13 develops a notion of visual *reference*. Here I make use of the two-visual-system thesis of David Milner and Melvyn Goodale. The main idea is that visually guided action is hierarchical. At the level of consciously chosen action, the kinds of visual classification that have been discussed above are operative. On the other hand, the implementation of such action requires motor control of limbs and muscles relative to external objects—the precise reaching out to objects, the choice and sizing of grip relative to them, and so on. These control mechanisms are under the control of visual systems that do not provide explicit details to visual consciousness—in *SDK* they are collectively entitled "motion-guiding vision."

Motion-guiding vision does not provide visual consciousness with data that feed into epistemic operations. Rather, it gives the limbs data that enables them to grasp and to manipulate objects: these data are relational and transitory, and in contrast to the information provided by sensory classification, not about persisting qualities of any object. When successful, motion-guiding vision puts us into a relationship to objects that involves these objects as objects of manipulation, not as the bearer of visual features. Thus, visual states result from data-processing of two quite different kinds: they comprehend visual features, which are attributed to objects apprehended independently, by being apt for direct physical manipulation. When I look at an object within reach, (a) my visual state takes in its sense-features, and (b) motion-guiding vision connects me directly with the object. The visual state has an object-directedness which derives from the ability to grasp and manipulate objects that are within reach, and a descriptive element that derives from sense-features being attributed to this object.

Seeing, Doing, and Knowing constructs a comprehensive framework within which we can understand the form and content of the information that the senses provide about the world outside. It shows how the senses link with objects that it presents under various sense features, and gives a new species-specific account of the essence of these features.[2]

[2] Warm thanks to Vincent Bergeron, Jonathan Cohen, Larry Hardin, Benj Hellie, Janice Lindsay, Dom Lopes, and Bence Nanay.