REPRESENTATION OF KNOWLEDGE: NON-LINGUISTIC FORMS

DO WE NEED IMAGES AND ANALOGUES?

Zenon W. Pylyshyn

Department of Psychology
University of Western Ontario
London, Canada

. OLD HOMONCULI NEVER DIE

It is no accident that inside most sychological theories of representation we an, if we look closely enough, discern a mall person with his eyes on a screen and is hands on the controls. The metaphor is o seductive that almost all theories of erception succumb to it (as Kaufman, 1974 as noted in his recent review of theories n perception). True, we try to deliver the omonculus a better and more stable picture han falls on the eye of the larger person e is controlling -- in fact we usually go o the trouble of presenting him with a hree-dimensional model (often holographic), oping to lighten his load, but the little an seems so friendly and familiar that we an't imagine how we could do without him. he dilemma this places us in goes back everal millenia. It runs something like his. We need to have some internal epresentation of the world in order to hink about it (indeed, in order to pprehend it at all). But if this internal epresentation is too similar to the world tself it cannot help us to apprehend it ince it merely moves the same problem nside. On the other hand if it is too issimilar then how can it represent the orld at all? Epistemologists have squirmed nder the horns of this dilemma trying by arious means to make the problem disappear. sychologists on the other hand have by and arge dismissed the problem as old-fashioned which it is) and have proceeded to be igorous in their experimental analysis of he "functional role of images", where mages are not merely "pictures" but are rtfully becoming much more fleeting and ketchy. Sometimes they are referred to as erceptual schemas, sometimes as "the ctivation of perceptual processes", and ore recently as analogues. The little man or his part has been put in a black box here he continues to live under such guises s "the visual system" or as something which esponds to the analogues by moving limbs or ttering sentences as required. This ccount is admittedly unfair to the many nvestigators who understand the basic roblem quite well and are struggling to evelop representation systems adequate to he task. But I believe that the caricature dequately characterizes the vast majority f psychological approaches to the henomenon of so-called "non-verbal epresentation".

I will confine my written remarks to a mall subset of questions bearing on this ilemma. I would be glad to provide eprints of my other relevant papers on equest. Primarily what I will try to do is to point out that many of the ways of casting the problem of "alternative forms of representation" are misguided and that by blurring certain distinctions and emphasizing others we may be burying the significant problems in a mire of catchwords (e.g., procedural embedding, analogical, holistic and even propositional -- which I now regret using because of the sentential connotations which, despite all my efforts, it continues to have).

II. THE FUNCTION OF REPRESENTING: "RESEMBLING" OR "DESCRIBING"?

Let us look at the representation dilemma again. It asked (in part) how an entity could represent some object if it was too dissimilar from that object. But this, like a great many other questions of this sort, already presupposes something crucial. We normally only speak about two things being similar if they are to be examined in the same way -- in particular if they are both to be viewed. Since we don't want to start off with this as the assumption (we might then ask "who does the viewing inside?") we should drop the idea that the representation literally resembles the thing it represents (see Goodman, 1968, for more on this point). Well then can the representation be any arbitrary symbol? Clearly it cannot in general be an unstructured atomic symbol since then there would be no way to show that the thing represented had a structure -- i.e., had subparts, relations and attributes. So what constraints are there on the structure of the representation? Here the going gets tougher. One is tempted to give the recursive reply that it must have substructures, relations, properties, etc. which represent the substructures, relations and properties of the object(s) being represented. But here we have to be careful for two reasons. One reason is that if the representation maps all the structures, etc. of the object we will have an isomorphism which has all the disadvantages of the picture-in-the-head alternative. The representation must not only be highly partial but it must be partial in the appropriate way (see below). The other reason is that it is meaningless to speak of the structure of the representation. Structure is relative to the processes which construct and use the representation. It is these processes which define the semantics of the representation: we may speak of the structure of a representation relative to a Semantic Interpretation Function (SIF). Thus the two distinct strings of symbols "not (p and q)" and "not-p or not-q" are identical structures from the point of view of a theorem prover and the distinct strings "* +" and "(LEFT-OF STAR PLUS)" may be identical structures from the point of view of some other SIF. Neglect of the SIF represents one of the most ubiquitous sources of confusion in discussions about representation. It leads some people, for example, to assert that non-linguistic representations "preserve the structure of that which they represent". They do so of course only to the extent that the "same

structure" is extracted by some appropriate SIF. In that sense the sentence "the book is on the table" can be said to preserve part of the structure of a scene containing a book on a table. To be sure the latter has a lot more structure as well but so does the sentence (it has order, length, color, etc.). It is up to the SIF to pick out those aspects which are signifying from those that are not and to process the string (in the appropriate contexts) as it would the scene. Without knowing what the SIF did we could not speak of structural similarity. I don't mean to imply by this example that sentences provide an adequate representation of scenes (they don't for other reasons) but only that the differences are more subtle than captured in the simple claim that the scene and the sentence have different structures. At this level all we can say is that they don't "resemble" one another.

One can of course remove much of the arbitrariness in the above characterization of the structure of representatons by requiring that the SIF be perceptual in nature -- i.e., by assuming that the SIF has much in common with visual perception. There is a good deal of psychological evidence suggesting that imaging and perceiving are similar in many ways. Although this seems like a reasonable proposal it creates many problems and must be approached with care. It is tempting to "explain" aspects of cognition (e.g., Moyer's (1973) account of magnitude judgments from memory) by pointing out that they are "like" their perceptual counterparts in respect to such measures as reaction time. But since we have no idea of how the latter is accomplished this is a case of "obscurum per obscurus". Furthermore to note that some cognitive operations bear a (not yet well understood) relation to perception is in no sense evidence that these cognitive operatons involve pictorial or analogical or any other entities which resemble objects in the environment. Presumably perception involves the construction and processing of internal representatons just as does imaging so some relations between the two should not be too surprising. Furthermore there are some major differences as well. These are related to the fact that objects in the environment have a stable existence so they can be re-examined and to the fact that transformations of internal objects (such as those studied by Shepard) depend on the person's tacit knowledge concerning permissible transformations. The way in which this knowledge must be brought to bear -- and not intrinsic properties of the representation (i.e., not the rigidity of patterns being mentally rotated) are what must account for experimental results on mental transformations (we shall return to this point in section III).

But perhaps the main argument against the view that the SIF is perceptual -- assuming that we can specify what we mean by perceptual in other than hand-waving terms -- is that it implies that the representation to which it is applied is something capable of being perceived. Unfortunately no matter how hard we try to make it sound like we are avoiding pictures (or worse, objects) in the head there is no coherent intermediate ground: if the SIF has perceptual primitives (e.g., operations such as those studied in vision for feature detection, etc.) it must be applied to something which, however fleeting, sketchy, vague, dynamic, etc. is still pictorial or isomorphic in a sense which is incompatible with the facts of human memory and cognition. I want to make it clear that I don't object to the reification of pictures or some such analogues on ontological grounds, but simply on the grounds that such objects as a class have the wrong properties. Our representations of the visual world are not like any (degraded, topologically transformed, filtered, etc.) projection of proximal stimulation: they are constructed from aspects of the world which we <u>notice</u> (and such aspects can be global, abstract and highly cognitive -- i.e., knowledge-driven and assimilated into available conceptual categories) and they represent equivalence-classes of stimuli which are physically very different from each other and from any conceivable picture-like entity. For example I might notice shapes (or at least a class of shapes) but not colors, objects but not locations, and non-sensory relations such as causality, potential actions, intensions, etc. Such representatons, derived from visual perception, cannot be sharply distinguished from knowledge derived by other means; that is why I prefer to refer to them as "structured descriptions". The vocabulary of such descriptions and the accessibility relations may be quite different from that of linearly ordered utterances. Such "visual images" are in some ways more like models than logical statements insofar as they may not contain quantifiers (at least the current computational models of imagery do not -- e.g., Baylor 1972, Moran 1973). Images in such an approach are data structures in which objects are individuated (i.e., there is no node for "seven blocks"), contain many "default" attributes and typically use spatial relations as access paths. Yet in my view it is more appropriate to refer to them as descriptions than images because the term is less misleading since they consist of conceptual structures very much like those constructed when the input is linguistic -- except perhaps using a modality-specific vocabulary of symbols. One cannot of course rule out the possibility that there are cognitively functional aspects of percepts which cannot be captured in such a discrete symbol system, but I have yet to hear a persuasive argument for that case. Furthermore, I have argued elsewhere (Pylyshyn, 1973) that there are many conceptual traps awaiting those who talk in terms of storing and using images.

III. ANALOGICAL AGAIN

The most common proposal for an alternative form of representation for perceptually derived knowledge is that it is analogical. This term has become the new

uzzword in cognitive psychology and is used s a synonym for anything from "warm and uddly" through "holistic", "continuous", or imply "anything which is not anguage-like". Few psychologists have ried to be very specific in characterizing he meaning of this term. When people have ried to be explicit (as, for example, loman 1971; Block and Fodor 1973; Lewis 971, Goodman 1968) they have found it to be very difficult concept to characterize and ave had to distinguish several different enses in which the term is used. I have iscussed some of these elsewhere (Pylyshyn, n press) so I will not repeat myself here. want merely to add to what I have written ome discussion of why people may be tempted o reach for analogues to account for ertain psychological evidence, and to uggest why such entities whatever they may e, fall short of serving the function xpected of them.

As a psychologist one of the main bjections that I have to the whole notion f analogue representation is that it seems o me to be a convenient way of hiding a arge part of the problem we are trying to xplain -- i.e., how people represent and eason about objects and actions. You may ecall being at least mildly surprised that here is such a thing as a "frame problem" n reasoning about actions (McCarthy and ayes, 1969; Simon, 1972). The reason that t never occurred to many of us that there as a problem is that when we interact with he environment (as opposed to thinking bout it) the laws of physics take care of ll the relevant interactions among vents -- we don't have to worry about verlooking what will happen to evrything lse in the world when we carry out some ction on a part of it. Such relations are iven to us free by the environment. In the ase of reasoning, however, the relations re not free. We must in some way xplicitly build in the knowledge regarding hat effects do and don't follow from any ction. Now it seems to me that the notion f an analogue representation is in part an ttempt to get this information for free gain. Thus the claim that data on the ime-course of mental rotation (c.f., Cooper nd Shepard, 1973) argues that the process s analogue (since, as the proponents nnocently ask "how can you rotate a data tructure through its intermediate ositions?"). This carries the implication hat once we start a rotation the medium ill take care of maintaining the rigidity f the total pattern and carry along all the arts for us -- just as the laws of physics ake care of this for us in the real nvironment. But, as in the frame problem, e are overlooking the fact that the person or the robot) must know what will and will ot happen to the bottom part when the top art starts to rotate. In a descriptive tructure this is precisely what makes 'mental rotation" appear awkward and omputationally unduly costly. But this is navoidable unless we have an analogical odelling medium which intrinsically follows he laws of physics. Unless we are willing o ascribe such laws to brain tissue (which, y the way, is what Gestalt psychologists attempted to do) we are stuck with locating it in what I have called the SIF (which does not, incidentally, preclude it from being a distributed computation attached to the data structure itself). If we admit this, however, there appears little reason to call the resulting representational system analogical (though Shepard's use of the term is, by his own admission, broad enough to cover this case).

Another example where analogues are invoked in a similar role is for the representation of magnitudes. When we "mentally compare" two objects -- say a dog and a horse -- to judge which is larger, the answer seems immediate and intuitively appears to depend on a comparison of two images or some sort of "analogues". Now we have some idea of what sort of operation is involved when two physical objects are compared by placing them side-by-side. Again the laws of physics and optics assure us that, as in the frame problem, the right things will happen (e.g., the object sizes will remain fixed as they are moved, the smaller object will partially occlude the larger, etc.). But in the mental comparison case we somehow feel that the analogues will "do the right thing" because of intrinsic properties of the analogue medium, just as in the mental rotation example we feel that analogues will intrinsically maintain their form in a rigid manner during rotation. In the mental comparison case the assumption is that if the process is analogue, the SIF does not need to "know" the rules of transformation nor does it need to "know" about order relations -- e.g., that such relations are asymmetric and transitive -- since it has merely to "read off" the answer from the analogue. The representation again seems to have the answer "written on its sleeve". Thus by attributing such properties to the intrinsic nature of the representation we beg the very question of how magnitudes are encoded and compared.

The phenomenon of attributing to the intrinsic nature of a representation some of the crucial aspects that need to be taken into account (because these are so intuitively obvious to the theorist) is not confined to analogical representations. Woods (1975) has recently shown that we frequently commit the same oversight in the case of semantic networks. This is why it is important to attempt to simulate a significant portion of cognition by machine (although even here the existence of such built-in functions as an arithmetic processor may create the illusion that we get magnitudes for free -- i.e., we need not model them in detail).

In conclusion let me reiterate that I don't claim to have made an argument against analogical modes of representation -- and still less that I am satisfied that semantic networks, procedures, etc. are adequate to handle all forms of knowledge. I have simply tried to argue that many of the reasons people have for jumping on the "non-linguistic" (whatever that may be) bandwagon are insufficient. Furthermore we

are so far from understanding the semantics of discrete data structures (as Woods has cogently argued) that any mass movement to abandon them (or even augment them with something radically different) is at the very least premature.

REFERENCES

Baylor, G.W., A treatise on the mind's eye: An empirical investigation of visual mental imagery. (Doctoral dissertation, Carnegie-Mellon University) Ann Arbor, Mich.: University Microfilms 1972. No. 72-13, 699.

Block, N.J., & Foder, J.A., Cognitivism and the analog/digital distinction, Mimeo, MIT, 1973.

Cooper, L.A., & Shepard, R.N., Chronometric studies of the rotation of mental images. In W.G. Chase (Ed.), Visual information processing. New York: Academic Press, 1973.

Kaufman, L., Sight and Mind, New York: Oxford University Press, 1974.

Lewis, D., Analog and digital. Nous, 1971, 321-327.

McCarthy, J. & Hayes, P., Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer & D. Michie (Eds.) Machine Intelligence 4, Edinburgh: University of Edinburgh Press, 1969.

Moran, T., The symbolic imagery hypothesis: a production system model. Unpublished Ph.D. dissertation, Carnegie-Mellon University, 1973.

Moyer, R.S., Comparing objects in memory: evidence suggesting an internal psychophysics. Perception & Psychophysics, 1973, 13, 180-184.

Pylyshyn, Z.W., What the mind's eye tells the mind's brain: a critique of mental imagery. Psychological Bulletin, 1973, 13, 1-24.

Pylyshyn, Z.W., The symbolic nature of mental representaions. In S. Kaneff and J.E. O'Callaghan (Eds.) Objectives and Methodologies in Artificial Intelligence. New York: Academic Press (in press).

Simon, H.A., On reasoning about actions. In H.A. Simon and L. Siklossy (Eds.) Representation and meaning. Englewood Cliffs, NJ: Prentice-Hall, 1972.

Woods, W., What's in a link: foundations for semantic networks. In D. Bobrow and A. Collins (Eds.), Representation and undrstanding: studies in cognitive science, New York: Academic Press, 1975.