

difference causes for quasi-pictorialism. It raises other problems for descriptionism, I believe, and we will soon turn to them. Before we do that we will try, in the next section, to get a more concrete sense of what is meant by "propositions" and "descriptions" in these contexts. Before we do even that, however, we should take note of at least one significant way in which neither of these theories scores over the other. The account of our conscious experience in both quasi-pictorialism and descriptionism is the same. In both cases consciousness is a matter of having propositions in STM. Whether they have got there from LTM directly or indirectly makes no difference; in either case they will be equally like the propositional end-products of perception itself. The visual processing activity postulated by Kosslyn's type of theory will be quite unconscious (just like the rest of visual processing) and so, perhaps surprisingly {28\*}, neither theory will give a better account of the immediate experience of imagery (or perception) than the other. The corollary of this is that both theories stand or fall together on the viability of this 'propositional' view of the mind. That issue will be deferred until the final section of this chapter.

### **§II.C.2. Simulating Imagery in Terms of Descriptions.**

In the early days of AI research in the U.S.A. there was a certain amount of amicable disagreement between

those who saw the problem in terms of getting computers to display 'intelligent' behaviour any way this could be done and those who were concerned with getting computers to mimic specifically human thought processes, with what they called "cognitive simulation". The former approach was particularly associated with Marvin Minsky of MIT, and the latter with Herbert Simon and Allan Newell of Carnegie-Mellon University {1\*}. One upshot of this difference is that the Carnegie-Mellon school and their successors have been much more open to influences from developments in psychology (both approaches have had significant influences on psychology). In particular, they felt it needful to take seriously the upsurge of interest in imagery amongst cognitive psychologists towards the end of the 1960s. By the early 1970s Simon and Newell were actively thinking about how cognitive processes involving imagery might be simulated on computers {2} and three of their students, George Baylor, Thomas Moran and Arthur Farley, attempted such simulations for their doctoral theses {3}. As previously noted, it is largely this work, particularly that of Baylor, which has encouraged some psychologists to believe that imagery can best be accounted for in terms of 'propositional' descriptions, and we shall therefore describe it in more detail here. The little known work of Farley, however, has distinctive features which will cause me to postpone its full consideration until a later chapter. It is also worth noting that since this work was done (i.e. since 1974) little or no further work specifically on imagery has been generated from within AI



{4\*}. I fear that this may reflect a failure within AI to confront the problem which, we have argued, the imagination was invented to explain - the problem of how rational thought gets in touch with empirical reality. But we must leave proper consideration of such matters until later.

It has been characteristic of all schools of AI to attempt to produce programs which mimic a very restricted range of mental abilities to a more or less fully human level of competence, rather than to try to produce 'whole minds' (even very stupid ones) in one go {5\*}. From its inception the Carnegie-Mellon school centrally concerned itself with simulations of human 'problem solving', that is the solving of puzzles such as those of "cryptarithmic" {6}, an example of which is given in figure II.C.2\_1. It would be relatively easy to program a computer to solve this by 'brute force' as it were, simply by trying out every number letter combination until one was found that fitted. However, doing things this way would be of no psychological or philosophical interest whatever. Human beings are very much slower at simple arithmetic than computers, and it is clear that such an approach is not practicable if people are to solve such problems. Simon estimates that it could take several years of full time work for a person to solve the problem of figure II.C.2\_1 by such a brute force method {7}. Since people often can solve such problems relatively quickly they must be using a more efficient strategy. Generally speaking, Newell and Simon believed, such strategies differed from 'brute force'

methods in not being algorithms, - i.e. in not being certain methods of finding the answer - but they generally had a very good chance of returning an answer fairly quickly. Newell and Simon called such strategies "heuristics", and they were interested in finding out just what are the heuristics which people actually use; both as a study in human psychology and as a step towards getting computers to 'think' like people {8\*}. Possible human heuristic strategies could be tested by programing them onto a computer and seeing whether, if they worked at all, the pattern of performance for person and machine was similar (e.g. problems which people find more difficult should take the computer concomitantly longer). One of the most important methods employed for generating hypotheses about the heuristics which people might be using was a technique known as "protocol analysis" {9}. This simply involves the human problem solver (often the experimenter himself) 'thinking out loud' as far as possible whilst he attempts to do the puzzle. These 'thoughts' are tape recorded and thence transcribed, and from both the content of the words and from the pattern of pauses (often from just a single example) hypotheses about the steps and strategies of problem solving can be derived.

What both Baylor and Moran did was to apply this methodology - protocol analysis followed by testing the hypothetical heuristic on a computer - to problems which subjectively seem to use mental imagery in their working out {10\*}. The underlying assumption seems to be that the

use of these heuristics, these particular patterns of information processing, amounts to imaging (or at the very least it directly causes images to be produced, as epiphenomena, when it occurs in brains) {11}. The problem task used by Moran was to follow a path in two-dimensional space. The subject was given a succession of cardinal compass directions (e.g. "north - east - south - east - south - west...") and was to imagine a point moving a fixed distance in each of these directions, tracing out a line, and had to report on the final position of the point relative to the start. The protocol produced gave indications of such things as the sorts of simple shapes into which the whole complex path was conceptually broken down (e.g. "north - east - south" might provoke the description: "a box open at the bottom"). These could then be used as representational units in the computer's program {12}.

The type of problem tasks used in the better known work of Baylor were the "Block Visualisation Tasks" devised by Guilford {13} to test spatial thinking ability. These tests require us to report on the consequences of variously painting and slicing up certain cuboidal blocks. Baylor recorded himself talking through the solving of the following task:

The four narrow sides of a 1 inch by 4 inch by 4 inch block are painted red; the top and bottom are painted blue. The block is then cut into sixteen one inch cubes. (1) How many cubes have both red and blue faces? (2) How many cubes have one red and two blue faces? (3) How many cubes have one red and two blue faces? {14}.

From consideration of the protocol thus produced (together with his knowledge of Paivio's "dual coding" theory of memory {15}, to be discussed later) Baylor concluded that two mental representations of the block are involved in solving the problem. The first of these, which he calls the S-Space representation (S=symbolic), seems to correspond to how we might describe features of the block to ourselves verbally. The second representation he calls the I-Space representation, and it is supposed to underlie the mental images which seem to be involved in doing such a problem. This I-Space representation is, nevertheless, conceived of as being of the form of a description, in the form of certain symbolic terms and their relations. Figure II.C.2\_2 is a list given by Baylor of some of the relevant symbols employed by his program.

S and I-Space, it should be noted, are only spaces in a highly metaphorical sense. They are not even meant to be different parts of the mind; rather they are "problem spaces" {16}, separate only inasmuch as they involve different sets of representational terms and relations on which different sorts of operations are carried out. Since both representations are in fact forms of symbolic description, drawing a sharp distinction between them may be rather misleading. Simon in fact tells us:

that all internal modalities employ basically the same kinds of relational structures for storage; and that their differences are differences in organisation which enable them to be accessed by processes paralleling the processes of the corresponding sensory modalities. {17}.

Where Baylor is inclined to stress these organisational differences, Moran, whose approach to the simulation of imagery is otherwise very similar to Baylor's (18\*), stresses the commonalities:

There is no separate image memory, or imaging sub-system. Visual information is just like any other information in the system: it exists as symbolic expressions in STM which are acted upon by rules in LTM. (...) [It] is hard to draw a clean line around the 'purely visual' operations - they are all mixed up with non-visual operations. (19).

Thus:

My model (...) has only one abstract 'semantic' system - a structurally homogeneous representational system. STM contains neither mental pictures nor mental sounds but only symbolic descriptions of them.

(20).

Be that as it may, it is those representations which Baylor segregates into his I-Space which particularly interest us. We want to see how these amount to 'descriptions' which nevertheless have a distinctive type of structure which differentiate them from the sorts of descriptions we give in ordinary language (or in S-Space symbols). The structure of the I-Space description is supposed to make **explicit** (and thus more accessible for problem solving purposes) those properties and relations which seem to be immediately available in visual experience but which would be merely **implicit** in an ordinary language or S-Space description. This might seem problematic because, like language, current computers deal with symbols sequentially. Of course, one could make a sequential

symbolic description as fully explicit and 'visual' as required simply by writing it as a very long conjunction:

A block has this face (#F1) and that face (#F2) and that face (#F3) and ... and face #F1 has this edge (#E1) and that edge (#E2) and ... and face #F1 has this vertex (#V1) and ... and face #F6 has this edge (#F3) and ... and this vertex (#V1) is to the left of that vertex (#V2) and .... {21\*}.



This however would be extremely unwieldy. We might have to work all through this list of properties before we found the one we wanted. When we imagine a block, by contrast, we seem to be able to go straight away to 'look at' whichever face we want, and from there to be able to 'see' its edges and corners. To 'see' the edges and corners of another face, however, we must first be 'looking at' this second face. We want our I-Space descriptions to have a structure which enables them to be accessed in this sort of visually ordered way {22}, and this can be done by making the representation hierarchical {23}. Thus Simon gives the following account of 'image' representation within one of his programs:

An image is simply a short list of symbols (two or three symbols), each of which is either a primitive symbol (not further analysable) or the name of some other image. Thus we might have: I1=<I2,I3>, I2=<I4,P1,P2>, and so on, where the I's are the names of compound images, or chunks, while the P's are the names of primitive symbols. {24}.

Simple list structures like this are not, Simon admits {25}, sufficient to account for the properties of mental imagery. Such a structure might be able to express, for example, that an image contained a man and a cigar (if MAN and CIGAR were primitives!), but it could not express

that the man was smoking the cigar. However, if instead of only being able to express simple conjunction the syntax of the representational system is enriched so that it can contain propositions expressing property possession (CIGAR ~~is-a~~light) and two termed relations between objects (MAN smoking CIGAR) then, Simon claims {26}, we will have sufficient representational power to properly account for imagery. MAN and CIGAR, of course, would not really be representational primitives. The idea is that the images of the man and the cigar are themselves decomposable into descriptions in terms of lower level images, probably through several levels, until we reach the true representational primitives. These primitives are the symbols which are supposed to represent certain simple features of the real world in an arbitrary, 'linguistic', way.

The full decomposition of a suitable 'I-Space' representation for an ordinary object like a man or a cigar to any plausible primitives would be horrendously complicated, and no-one, I think, has attempted to show how it might go. However, if we consider simple geometrical figures and solids (and neither Baylor's, Moran's nor Farley's programs attempt to deal with anything more) the basic idea can be illustrated. Simon {27} attempts this with a simple straight line ABC, but the later account given by Palmer (with acknowledgement to Baylor's work) is perhaps more clear [Palmer uses **bold** type for terms denoting symbols within the representation, and also



The structural description of square ABCD. Nodes represent the component POINTs, lines, and angles of the square (see Figure II.1A). Parameters of the component parts are not shown.

Figure II.C.2\_3  
(Reproduced from Palmer [1975].)

**CAPITALISES** the terms denoting the primitive symbols]:

The structural units can be primitive (for example **POINT**) or higher order (for example **squares** and **cubes**). In a propositional system, an **angle** might be represented by a proposition expressing the relation between the component **lines** and the **vertex** point of the angle. Similarly, a **triangular pyramid** might be represented by a set of propositions about the intersection between the four component **triangles** at the six **edges** and four **vertices** of the pyramid.



{28}.

The structure of such representations can become quite complex even for relatively simple perceptual shapes, as Palmer's examples show:

Figure [II.C.2\_3] illustrates the complete {29\*} structural description of **square ABCD**. The figure shows the structural **PART** relationships between the square's component parts (**lines** and **angles**) and the parts of its component parts (**POINTS**).

{30}.

To explicate this diagram somewhat, all object terms are represented by symbols in angle brackets, whilst relational terms are represented by the labelled arrows. Thus each relational proposition is shown as two object terms joined by an arrow. The square as a whole is represented by the symbol **<ABCD>** at the central node. To be holding or manipulating the equivalent internal symbol in STM would be to think about a square as a whole, a square gestalt {31\*}. The square can either be immediately analysed into sides **<AB>**, **<BC>**, **<CD>** and **<AD>**, or into angles **<ABC>**, **<BCD>**, **<CDA>** and **<DAB>**. Angle **<ABC>**, for example, is stated to be part of **<ABCD>** and has as its own parts lines **<AB>** and **<BC>** and vertex point **<B>**. Lines such as **<AB>** are also directly stated to be parts of **<ABCD>** and themselves to be composed of point parts such as **<A>** and **<B>**. The idea is that when



The structural description for cube ABCDEFGH expanded to the level of lines. (Vertices are not shown.) Nodes representing parts of the cube not visible from the perspective shown above are within the shaded area.

Figure II.C.2\_4  
 (Reproduced from Palmer [1975].)

we visually or imaginatively consider the structure of a square we follow the arrows. We may be led from seeing or 'seeing' the square as a whole (having  $\langle ABCD \rangle$  in STM), via the proposition that  $\langle ABC \rangle$  is part of  $\langle ABCD \rangle$ , to imaging or mentally fixating one corner ( $\langle ABC \rangle$  is now in STM); and thence in a similar way to imaging one side ( $\langle AB \rangle$  or  $\langle BC \rangle$ ) or one vertex ( $\langle B \rangle$ ). What one cannot do is move directly from  $\langle ABCD \rangle$  to one of the point symbols, or, for example, from  $\langle ABC \rangle$  to  $\langle AD \rangle$ . Thus, the claim is, the order in which such structural descriptions must be 'read' accounts for the order in which we seem to 'read' images or actual visual objects.

Palmer gives a similar diagram (figure II.C.2\_4) for the structural description of a cube (but decomposed only to the level of edges). This is closely akin to the representation structure for a block actually used by Baylor. Palmer's representations do not seem to be entirely adequate to imagery (they are not presented as such) because they deal only with the object's internal structure and leave quite indeterminate matters such as its spatial orientation, which are surely quite definite in an image (even Fodor's stick tiger is surely facing a particular way). Baylor's actual block representation, which is presented as a model for imagery, includes such orientation information:

The full image or I-Space representation of a piece or block is a list of eight vertices, called a piecelist; this piecelist is connected to six facelists (each a list of four vertices), each of which is connected to four edgelists (of two

