14

An extension of the theory to humans

14.1. Introduction

Any model of hippocampal function must account not only for data from the rat, cat, and monkey, it must make sense of data gathered in the clinic as well. For over 20 years it has been known that the hippocampus is essential in human long-term memory function, and the inability of investigators to find adequate parallels to this function in other species led to the unfortunate tendency to assume that there were major species differences regarding hippocampal function. With the advantage of hindsight we can see that this discrepancy between clinical and experimental work resulted from two related factors: (1) an inaccurate conceptualization of the nature of the memory deficit consequent upon hippocampal damage in humans, embedded within a similar set of misconceptions concerning the nature of normal human memory; (2) the use of inappropriate paradigms in work with infra-humans, which generally failed to assess the same aspects of memory affected by hippocampal damage in humans.

We have argued elsewhere (Nadel and O'Keefe 1974) that there is no valid reason to assume that differences of large magnitude exist between mammalian species in regard to the function of specific brain areas, and we are thus in strong agreement with the recent tendency to bring the clinical and experimental work together under the same theoretical umbrella (cf. Isaacson and Pribram 1975).

In the first section of the book we established the basis for the theory of cognitive mapping by considering data from a variety of species, including humans. Impressive evidence for the existence of spatial mapping comes from recent work with humans, while the philosophical background to the problem of space perception rests almost entirely on considerations of the proper basis for human knowledge. The cognitive map in infra-humans should be viewed as a spatial map in which representations of objects experienced in the environment are ordered within a framework generating a unitary space. However, the central property of the locale system is its ability to order representations in a structured context. The development of objective spatial representations is not the only possible use for such a system. We hope to show, in this final section of the book, that mapping structures can represent verbal, as well as non-verbal, information. For both of these forms the locale system will be shown to be central to a particular form of memory: that concerned with the representation of experiences within a specific context. We shall argue that memory comes in two basic varieties: (1) memory for items, independent of the time or place of their occurrence; (2) memory for items or events within a spatio-temporal context.*

In this first chapter of the final part of the book we shall consider two theoretical issues. First, we must consider the problem of spatial orientation in humans in more detail, in order to define precisely the nature of the spatial representation supposedly subserved by the locale system. This is necessary primarily because there exists considerable confusion in clinical neuropsychology concerning the various forms of spatial representation and their neural bases. Much of this confusion results from the empiricist bias in theories of space perception, as we noted earlier (p. 52). Second, we shall discuss, in broad outline, the general problem of human memory. Amongst the reasons which make it possible to draw together the clinical and animal experimental research is that theoretical conceptions of human memory have changed radically in recent years. In discussing memory, we shall touch upon the various forms it takes, and the need to see long-term memory, for certain kinds of information, in terms of holistic structures.

14.2. Neural correlates of human spatial representation

We have contrasted two forms of space perception: absolute, or non-egocentric, and relative, or egocentric. The former was seen as a unitary, objective space in which the position of the organism did not affect the distribution of objects represented within that space. The latter, on the other hand, consisted in a variety of spaces defined by the relation between the organism and external objects, or between different parts of the organism. The spatial mapping function proposed for the hippocampus is concerned with the former, non-egocentric, space; though it receives inputs from systems concerned with relative spaces, the hippocampus does not store these egocentric representations.

According to traditional empiricist psychology objective space is generated from relative space, and this position has been incorporated into much of the clinical work in the form of the assumption that the same brain areas are responsible for both. Thus, Teuber and his colleagues

* This distinction is similar to that drawn by Tulving (1972) between *semantic* and *episodic* memory. Many other writers have agreed that there is an important difference between memory which is essentially context free and that which incorporates, and is defined by, context. There are important differences in the detailed specification of these various dual-memory systems and we do not want to identify ourselves with any particular version of the theory at present. The dichotomy between route structures and map structures, as elaborated in Chapter 2, captures the essential differences between the context-free and context-dependent memory. As we shall see shortly (p. 398), much recent work on human memory has concentrated on those forms of memory storage which go beyond the restricted possibilities of association-based list, or route, structures.

(e.g. Teuber 1963, Semmes et al. 1955, 1963) assume that both forms of space, which they label intra- and extra-personal space, are subserved by the parietal area; this assumption seems to be unwarranted.*

In a review of the broad area of spatial disorders in clinical patients, Benton (1969) noted that

> 'Most patients who show defective absolute or relative localization of stimuli in external space do *not* show concomitant disturbances in more general topographical orientation or topographical memory ... This dissociation testifies to the essential independence of those two broad types of disorder' (p. 217).**

The fact that localization of objects in space relative to the position of the observer can be dissociated from topographical, or geographical, orientation attests to the independence of egocentric and non-egocentric space, that is, the separation between relative space and spatial maps.

While Teuber (1963) argued that disorders of both forms can be seen after parietal damage, more recent evidence suggests instead that the parietal area is involved with purely egocentric space, and the related process of recognizing objects in different perspectives (e.g. Butters and Barton 1970, Ratcliff and Newcombe 1973). Damage to the parietal area yields perceptual disturbances as well, and these, according to Benton, are not correlated with disorders of geographic memory.

From our point of view it is important to stress that the parietal area is not central to spatial mapping for two reasons: First, the fact that there is extensive post-natal development of the parietal area has been taken as evidence in support of the view that objective space cannot become available in humans until as many as 6-7 years after birth. This position, acceptable to Piaget, fails to account for the facts of spatial perception and development, as we have already seen (pp. 78-9). Second, there is a strong correlation between disorders of geographic orientation and general memory (e.g. Benson, Gardner, and Meadows 1976, Van Buren and Borke 1972); this makes little sense if one places the former in the parietal region and the latter in the hippocampus.

Thus, we allocate to the hippocampus only that aspect of spatial representation concerned with non-egocentric space. We would not expect to see, in association with hippocampal damage, deficits in body schemas,

* Benton (1969) stated that 'there has been much theorizing to the effect that spatial knowledge of one's own body (particularly right-left differentiation) underlies the development and maintenance of the apprehension of spatial relationships among objects or events in the external environment and therefore that there must be a necessary and intimate relationship between disorders of the body schema and disorders of spatial orientation ... these theories make pleasant reading but they are not altogether in accord with the facts of clinical observation' (p. 223). Semmes *et al*. claim to have found just such an association, between two-point tactual discrimination and spatial orientation, but their claim that the latter disorder was multi-modal, and therefore a general one, is not supported by the data they present.

** Benton defines 'absolute' localization with respect to the body axis; thus it is part of what we term egocentric space. It is not to be confused with the notion of absolute space.

or left-right differentiation, or localization of objects relative to the organism. As in other animals, where a specific deficit in place learning was postulated, we suggest only that humans with hippocampal damage would be deficient in the representation and use of total, geographic environments.

14.3. Human memory

14.3.1. INTRODUCTION

Any item entering the human memory system gets routed through a series of stores, some organized hierarchically, others in parallel, each analysing it and storing it in a particular form. Some of these stores form what we have called the taxon systems; others form the locale system. We cannot provide an extensive review of the literature pertaining to these processes; the interested reader is referred to many recent books for further discussion.* Rather, we shall concentrate upon those areas of central importance to an understanding of the amnesic syndrome consequent upon hippocampal damage. We shall do this by first describing the properties of human memory within the taxon and locale systems, and then attempting to bring out the correspondence between these properties and those specified by other approaches to memory.

14.3.1(a). Memory in the taxon system. As we noted earlier, there are several stages of memory within the taxon systems. The initial effect of the activation of any representation within the taxon systems is a lowered threshold in the neurones coded for that representation. During this phase some kind of temporal-order information can be preserved in the form of the relative strengths of the traces of recently activated representations. However, since the long-term effect of such activations summate, so that the most frequently experienced items have the lowest thresholds, order information is lost within the long-term taxon memory system. This is another way of saying that storage within the taxon systems cannot take temporal context into account.

Taxon memory stores are governed primarily by the principle of *category inclusion.* Each taxon category is interested in, or coded for, different aspects of items, and any given item is potentially specifiable in terms of the set of taxon categories it activates. Similar items activate similar sets, sharing many of the same neural elements. The extent of neural overlap in these representations is a measure of the similarity between the items represented; it also provides the basis for understanding the *interference*

* Some good books are Neisser (1967), Norman (1969, 1970), Tulving and Donaldson (1972), Steinberg and Jacobovits (1971), Anderson and Bower (1973), Melton and Martin (1972), Norman and Rumelhart (1975), Weimer and Palermo (1974), and Bobrow and Collins (1975).

seen between related items, when context is not taken into account.*

The entry into taxon storage of several similar items results in maximum threshold reductions for those features which all items hold in common, with progressively less change in threshold for the less popular features. This gives rise to the *abstraction* of a *prototype,* which may be different from any specific item experienced but which nevertheless captures the central tendency of the group of items involved (e.g. Reed 1972, Rosch 1973, 1975, Smith, Shoben, and Rips 1974).**

Thus, taxon memory is characterized by several basic properties: (1) categorization by feature similarity; (2) the abstraction of prototypes representing the central features of the category members;[I] (3) the absence of spatio-temporal context coding; (4) the consequent absence of order, time and place information;[II] (5) incremental storage; (6) substantial decay with time.

14.3.1(b). Memory in the locale system. The locale system can be pictured as a long map extending from the past into the future, segments of which contain groups of representations in a fixed spatial (and hence temporal) relationship. These representations can be drawn from any of the taxon category stores and might consist, for example, of objects forming a spatial map or, of semantic categories, yielding an idea (see below, pp. 401-10).

The properties of storage within the locale system have already been enumerated; (1) preservation of spatio-temporal context; (2) single occurrence storage; (3) minimal interference between different representations of the same item; (4) multiple channels of access for the retrieval of any, or all, of the relationships embodied in the map. These properties, as we have seen, convey to the map features lacking in taxon stores, such as a freedom from dependence upon specific input or output modes. Thus, any of the relationships stored in the map can be retrieved from it by activating any portion of the map, whether or not these relationships were

* There are taxon stores based on various aspects of input items: physical form, meaning, etc. Some differentiation between items similar in one respect can be achieved by reference to a different storage system, where they might not be similar at all. One of the advantages of multi-attribute coding (cf. Underwood 1969) lies in this possibility. It appears that the activation of the different taxon stores occurs in rapid sequence, and that any representation, even the most primitive, can be maintained through active attention (cf. Posner and Warren 1972); we shall discuss these processes in the context of an analysis of short-term memory (pp. 386-9).

** These categories contain information about the extent of similarity between the members and about the relationship between each member and the prototype; the latter relationship is asymmetrical. That is, the `distance' from the prototype to any member is different than from that member to the prototype. These asymmetries are reminiscent of what is seen in subjectively loaded mental maps (p. 77) and indicate that, for both verbal and non-verbal information, the overlap model we have postulated is somewhat oversimplified. It is interesting to note that Leibniz's model, which does not seem to work for the locale system, could be applicable to the taxon systems.

[I] As Palmer (1975) pointed out, such central tendency prototypes are not useful in the description of all categories and other principles might be involved in the establishment, for instance, of a category concerned with items which do not vary continuously over a range of possible values.

[II] We are referring here to long-term information.

noticed at the time of input. This follows from the fact that the members of any map segment form the context for each item within that segment.*

An important difference between the rat and human locale systems lies in the development in the latter of optional *strategies* for placing items in the map, manipulating those that are already there, and relating different segments of the map to each other. This development is unlikely to have been precipitous, in the sense of the sudden emergence of language, but rather to have evolved gradually, with transitional stages seen in the cat and monkey.** In humans, items in the map can be moved from one location to another, generating new configurations which provide information about possible actions. The new feature provided by evolution, which possibly occurs in the cat, is the dissociation of this predictive mode from actual movements.

Further, items, and sets of items, not related to the present sensory array can be called up and manipulated. Thus, items which are never related in reality can be juxtaposed, substituted for each other, or related in other ways to create new patterns. This possibility, in addition to others allowing for the comparison of different map segments, provides the basis for operations which might be central to some of the mental activities usually referred to as thinking, imagination, and creativity. Finally, the existence of such optional strategies for information manipulation within the map, aside from enabling humans to imagine things they have never experienced (such as an elephant in a tree), lay the basis for a liberation of the mapping system from its connection with space; the relationship between items in a map need no longer represent an inevitable `real-world' spatial relationship. This feature will become obvious when we discuss the way in which the map `spatially' represents semantic relationships (pp. 401-10).

14.3.2. A MODEL OF HUMAN MEMORY

While we cannot discuss the entire range of work on human memory, or even a fraction of it, we can broadly outline the main features of memory. The distinctions we have just drawn between the memory processes

* The similarity between our map formulation and other recent proposals concerning long-term memory (e.g. Norman and Rumelhart 1975) will be brought out more clearly in later discussion (p. 398).

** We suspect that this is mirrored by changes in the role of the frontal cortex, portions of which exert more and more direct control over the hippocampus with phylogenetic development, culminating in direct fronto-hippocampal connections in the monkey (see p. 131). There is, as yet, insufficient information to extend our analysis, though it seems likely that the middle portion of the sulcus principalis is the crucial area (cf. Van Hoesen, Pandya, and Butters 1972). This area projects to the septum in some species and indirectly to the entorhinal cortex or hippocampus in others. Another indication of this shift is seen in the correlates of theta activity, which serves as the locative input (or coding strategy) in the mapping system. Correlated with movement in the rat in most instances, theta becomes more and more dependent upon non-movement processes in higher species, as seen in the shift towards a predominance of low-frequency theta. Finally, the increasing importance of the CA1 field (see p. 106), as documented by Stephan (1975), also attests to the shift in capabilities of the hippocampus with phylogeny. This field is responsible for the misplace system driving curiosity and information-seeking behaviour.

mediated within the taxon and locale systems will be presented in a more familiar context in what follows.

Norman (1970) summarized what can be taken as a representative view of the various forms of human memory:

'First, newly presented information would appear to be transformed by the sensory system into its physiological representation . . . and this representation is stored briefly in a sensory information storage system. Following this sensory storage, the presented material is identified and encoded into a new format and retained temporarily in a different storage system, usually called short-term memory. Then, if extra attention is paid to the material, or if it is rehearsed frequently enough, or if it gets properly organized, then information is transferred to a more permanent memory system' (p. 2)

though it is considerably out of date in some respects. Three sequential stages are postulated in this type of model: immediate memory, short-term memory, and long-term memory. Most traditional memory models, including those based on clinical observation, have emphasized the way in which inputs flow through a variety of passive information processors. More recent models, on the other hand, have moved away from

'rigid boxes of stores and towards a flexible system with a large part played by such optional control processes as attention, rehearsal and strategies' (Craik 1971, p. 236).

This shift reflects a profound change in thinking about the nature of long-term memory systems, accompanied by an interest in the active 'control' processes required for gaining access to, and effecting retrieval from, these systems. The distinction we have drawn, along with many others, between context-free taxon category memory and context-dependent locale memory is but one manifestation of this change.

Traditional associationistic models of human memory failed in much the same way that S-R explanations of animal behaviour failed.* Two related areas of research, concerned with the process of imagery and the basis for language comprehension, have combined to lay the basis for new conceptions of human memory. Before describing these developments we can briefly outline the properties of the earlier stages of memory referred to by Norman. These, it will be seen, correspond largely to what we have labelled taxon memory.

14.3.2(a). Short-term memory processes. The first physiological representation of an input, usually referred to as iconic (visual) or echoic (auditory) memory, can establish a persistence which outlasts the duration of the stimulus by approximately 150 ms if actual stimulus duration is brief. Extensive analysis takes place during this time, resulting in a list of

* Anderson and Bower (1973) recapitulate this history and propose a neo-associationist theory which improves upon, but does not ultimately save, this approach.

physical features which define the input; this presumably occurs in neurones like those described by Hubel and Wiesel (1962) and numerous other workers. This processing occurs automatically, unaffected by prior learning.* Interactions between separate inputs can take place, on the basis of physical (spatial and/or temporal) overlap, such that adjacent items affect one another's processing, as seen in studies of masking and metacontrast (cf. Weisstein 1969). The information provided by iconic and echoic processing allows the observer to *construct* the environment. First, the total stimulating field is segmented into coherent units as in figure-ground separation or stream segregation (e.g. Bregman and Campbell 1971). This segmentation precedes identification; evidence from cataract-removal patients (von Senden 1932) and from clinical studies (Macrae and Trolle 1957) supports this distinction between what Hebb (1949) called *figural unity* and *figural identity.* Second, a set of transitional stages leads to the development of identifiable *physical traces* (e.g. Crowder and Morton 1969) which could persist for as long as 10 s (Phillips and Baddeley 1971, Entus and Bindra 1970).** At more or less the same time, the 'meaning' of the input is encoded; that is, the neural representation within the *semantic* category to which items belong is activated. Thus, a word is coded both for its physical appearance (and/or acoustic properties) and its meaning within a brief period after input (cf. Posner and Warren 1972).

All these 'traces' constitute short-term activations within taxon category stores,[I] and the subsequent fate of the information stored in this way depends entirely upon the active control processes alluded to above. Though activation of a category representation will have long-term consequences, it is not sufficient to preserve spatio-temporal contextual information. This latter requires the establishment of memory within a different system; most control processes can be viewed as potential mediators of this mechanism.[II]

* We are excluding the developmental effects described by Hirsch and Spinelli (1971), Blakemore and Cooper (1970), and others. Should these prove to be true, they would be applicable only to an early critical stage in the development of the organism.

** The persistence elicited in these trace systems presumably underlies such phenomena as `priming' (e.g. Leeper 1935, Epstein and Rock 1960, and others), where the presentation of an item raises, for a brief period, the likelihood that it will be perceived upon subsequent presentation.

[I] We shall not discuss the evidence concerning hemispheric specialization at any length. Work with normal subjects, patients with unilateral brain damage, and patients with hemispheric disconnection has demonstrated that in humans the two hemispheres are largely specialized for different functions (cf. White 1969, Milner 1971, Gazzaniga 1970). Thus, taxon stores for verbal information would be located in the left hemisphere and those for non-verbal information in the right hemisphere. This specialization is preserved when one moves to locale memory, as we shall see later (p. 424).

[II] A number of memory-improving strategies depend upon manipulating inputs in ways which reduce the load on these context-free stores, though they still fail to preserve context in the spatiotemporal sense. Thus, nonsense syllables can be recoded as words (BLK-BLACK), acronyms can be used (International Workers of the World-IWW), individual words can be combined into a chain (natural language mediation, Prytulak 1971), and complex reduction rules can be applied to material as well as various substitution codes used, as described in the popular literature (e.g. Lorayne 1968). Such procedures can serve to decrease interference, but fail to achieve context-specific memory. In

The short-term memory resulting from activation of category representations has the characteristics of limited capacity, short duration, and the preservation of item and order information. This applies to both verbal and non-verbal information; however, there appears to be an important difference between the two. Aside from the various stores noted above, verbal information is apparently represented in a separate short-term system coded for the articulatory features of words.* This system possibly functions in both the comprehension and generation of speech, as we shall see later (p. 401).

Information can be maintained in the short-term memory state through the active process of rehearsal, which amounts to a recirculation of information and which might be restricted primarily to verbal items and to the articulatory code just mentioned. Rehearsal seems to be necessitated by the limiting features of the short-term memory process. The high reliability of retention for recent inputs and the strong order information available for such inputs indicates that access is achieved simply through 'testing' memory locations for recently used representations; those most recently used will have lowered thresholds. The capacity of the short-term memory system, in this view, is a joint function of the signal-noise detection capability of the system, the decay rate of the threshold change, and the rate at which representations can be activated. Rehearsal counteracts decay, but cannot enlarge the capacity of the system.**

Aside from its role in short-term retention and any specialized function in language comprehension and generation, short-term memory functions to maintain items in a state of accessibility until they can be encoded into a contextual long-term memory system. The set of items in short-term memory, which could include items called up from prior experience, must be encoded into long-term memory in such a way that activation of any member of the set activates retrieval of the remainder; in our view this involves locating the set in the locale system's cognitive map structure.

---

the same way the property of category inclusion can be utilized to improve recall, as shown in the work of Mandler (1967), Tulving (1966), and Bower *et al.* (1969), where knowledge of the categorial relationship between items in a list can serve to reduce the amount of information that needs to be retrieved. By themselves however, these reduction procedures would probably have minimal effects, unless they were made available at the time of recall (see p. 417). Their usefulness likely relates to the fact that they also reduce the amount of material that must be encoded into a contextual memory store, a process which requires time, and would therefore benefit from any such reduction in situations where speed of encoding is important. In contrast to these reduction strategies, techniques which facilitate contextual coding, as we shall see shortly, are useful when presented at the time of information input and are considerably more effective.

* Warrington and Shallice (1969) and Shallice and Warrington (1970) have provided strong evidence that this is a separate short-term memory store, located in a discrete neural area (the posterior parietal region), and not part of any causal short-term memory to long-term memory chain. The clinical syndrome of *conduction aphasia* represents the selective loss of this capability, with preserved long-term memory capabilities.

** Continued rehearsal will yield a 'better' long-term trace in the taxon system; the marginal effectiveness of pure rote learning is presumably due to the potentiating effects of sheer repetition.

We assume that finding and applying an appropriate coding strategy for a set of items takes time, and that short-term memory and rehearsal provide this time by maintaining the items in an accessible state.

*14.3.2(b). Contextual long-term memory.** Two areas of research attest to the central importance of context in long-term memory: (1) work on imagery (e.g. Bower 1970) indicates that storing information in an interactive, relational form vastly improves retention; (2) research in psycholinguistics indicates that *(a)* long-term memory for connected discourse (sentences, narrative) takes the form of retention of the meaning, or sense, of the discourse, typically with a loss of specific lexical elements and syntactic information (e.g. Johnson-Laird 1970), *(b)* analysis of the solely surface elements of language will not suffice to explain either the infinite generative possibilities of language or its comprehension, and (*c*) it appears necessary to postulate some semantic *deep structure* which incorporates the meaning of discourse without preserving lexical elements. Most contemporary investigators are in agreement in assuming that such a deep structure necessitates an abstract interactive system characterizing the relations between stored elements (cf. Weimer and Palermo 1974, Norman and Rumelhart 1975).

We would argue that a map structure can provide the basis for understanding the above facts and that it could serve as the form within which deep structure is articulated. Some discussion of work on imagery and language will support this position.

The function and nature of imagery. The demonstration that concrete words are recalled better than abstract words (cf. Paivio 1971) gave credence to the possibility that the chance to encode memory in imaginal form would improve long-term memory.** The independence of some form of imaginal coding from verbal coding had already been established (e.g. Brooks 1968) and the study of imagery was on its way. Subsequent work has verified the facilitatory effects of forming images between paired associates (Bugelski, Kidd, and Segmen 1968),*** the importance of an interaction, preferably spatial, within the image (e.g. Bower 1970), and the low interference seen between items remembered through the use of imagery coding.

The importance of spatial context in these effects is brought out clearly in the study of 'the method of loci', an imaginal technique known

* By contextual memory we mean coding within a spatio-temporal framework.

** Miller, Galanter, and Pribram (1960) had previously described the facilitating effects of a memory strategy dependent upon imagery, the one-bun, two-shoe technique.

*** Facilitation in this study depended upon a certain minimal inter-pair interval; this presumably relates to the time needed to find an appropriate image. This result has been replicated in other work.

to the ancients and described by Yates (1966) in her book *The art of memory* as well as by Luria (1969). In this technique the subject memorizes the layout of some building, or the arrangement of shops on a street, or any geographical entity which is composed of a number of discrete loci. When desiring to remember a set of items the subject literally 'walks' through these loci and commits an item to each one by forming an image between the item and any distinguishing feature of that locus. Retrieval of items is achieved by 'walking' through the loci, allowing the latter to activate the desired items. The efficacy of this technique has been well established (Ross and Lawrence 1968, Crovitz 1969, 1971, Briggs, Hawkins and Crovitz 1970, Lea 1975), as is the minimal interference seen with its use.

Results such as these had the initial effect of misleading some investigators into an oversimplified view of the nature of imagery (cf. Neisser 1972), epitomized perhaps by Paivio's assumption that imagery was quite distinct from language, that it involved a different form of coding than did the latter, and that it worked through the storage of something like a picture of the items to be remembered (cf. Paivio 1971). This latter assumption has been strongly criticized by Pylyshyn (1973), who suggested instead that some abstract structure generated in the process of imaging is stored in long-term memory, that this structure represents, though not literally, the relations between the elements in the image, and that the image must be reconstructed upon recall. Work with both visual images (Neisser and Kerr 1973, Kosslyn 1975) and sentence memory (Kosslyn and Bower 1974, Begg 1971, Anderson 1972) supports this constructivist view, indicating that the facilitative effects of imagery are not due to the long-term storage of a literal image, but rather to the storage of some deep structure analogue which can re-establish the image if required. In view of this, the facilitative effects of the 'method of loci' can be seen to derive from the ease with which the overall spatial image can be reconstructed. In addition, these newer data and the interpretation put upon them by most investigators seem to invalidate Paivio's notion that language and imagery work on different coding structures; both apparently demand some abstract deep-structure representation.

To recapitulate, imagery seems to work through the long-term storage of an abstract information set which captures the relationships between an interacting group of elements and from which the image can be easily reconstructed. This reconstruction process can start anywhere in the 'scene' and the image itself can be scanned in any direction (Lea 1975). This is consistent with the notion that the deep-structure analogue stored

* Kosslyn's (1975) study indicates that the literal image is reconstructed within the taxon system.

in long-term memory is one in which all the elements are interrelated and simultaneously available,* e.g. something like a map.

14.3.2(c). Language. The assumption that some abstract deep-structure base characterizes long-term memory first arose within psycholinguistics. We can briefly review this development, concentrating on evidence demonstrating that there are two stages in language processing, one involving the serial ordering of a linear string of symbols, the other consisting of an underlying non-linear structure from which this ordered string is generated. The similarities between these processes, and those just described for imagery, will become clear in the course of our discussion. As there is considerably more information available concerning deep structure in language than there is for imagery, we can go into some more detail here; the features of deep structure demonstrated for language are probably also to be found in the imagery system.

The existence of a deep structure. One approach to language, as to any form of behaviour, is to attempt to explain it solely on the basis of its observable linear structure. This approach, associated with the behaviourist school, suggests that each element of a sentence is generated in response to preceding elements, or in response to a stimulus in the environment, and that the whole sentence can be thought of as a Markov chain. Given the first word of a sentence, any other word has a finite probability of being produced, depending upon the number of times in the past that the word followed the first. Thus, a word like `smelly' would be followed quite often by `feet' or `cheese', less frequently by `music' or `airplane', and virtually never by `for' or `thinks'. Higher-order Markov chains would take into account not just the previous word but the previous two words, three words, and so on. Language, on this model, is generated solely by a system which produces strings of symbols in an ordered left-to-right linear sequence. Highly practiced sequences would be run off without recourse to decisions, ideas, etc. From this point of view there is nothing unusual about Lashley's colleague who claimed that '... he could arise before an audience, turn his mouth loose and go to sleep. He believed in the peripheral chain theory of language' (Beach *et al.* 1960, pp. 510-11).

Another aspect of language emphasized by the behaviourists is the

* There was some evidence in Lea (1975) that certain starting sites in a spatial array are more easily accessed than others; this seemed related to either top-to-bottom scanning methods or to some subjective impact of the objects located at those sites. This is, as Lea points out, an area requiring further research. Lea also failed to find any relationship between the reaction time required to scan from one site to another, and the real-world distance captured by the image of those sites; this is in disagreement with results reported by Kosslyn (1973). It is not necessarily the case, however, that a failure to find an increase in reaction times indicates that the image does not represent increased distances. There is no reason to assume, within a neural mapping structure, that real-world distances would be correlated with neural distances in a fashion which would produce orderly changes in reaction times with changes in imaged distances.

referential nature of meaning; that is, the way words refer to things or events and appear to derive their meanings from this reference. This connection comes about as the result of a simple conditioning process; a sound experienced in the presence of an object will, when later heard by itself, call up the same, or some of the same, responses as the object itself. The meaning of a sequence of sounds or words would then be given by the sum total of the conditioned meanings of each individual sound or word. In the face of harsh criticisms, to be mentioned below, this strong position has been progressively modified and weakened. In one recent formulation (Osgood 1971) meaning was seen as dependent upon some sort of internal response (r_m) which was derived from the total external response to the object. Words are not conditioned to the external responses but to these r_ms. The meanings of more abstract words, such as justice, are derived in a secondary fashion from the r_ms associated with actual objects or events.

This simple behavourist approach, which emphasizes the observable aspects of language, does seem to explain adequately many of the stereotyped features of language and some of the simpler referential features of meaning. It fails, however, in the language sphere in exactly those places where it fails in its explanation of behaviour in general, human or infra-human; it ignores or denies the purposeful variability and originality of behaviour, the novel behaviour not obviously due to generalization, the flexible use of behaviour learned in one situation but applied for a different purpose in another, and the underlying similarity amongst superficially different behaviours. These aspects of behaviour become acutely obvious in language and it is here that the deficiencies of the behaviourist account are most glaring. As pointed out by Lashley (1951), Chomsky (1957 and elsewhere), and Fodor (1965):

(1) Novel sentences constitute a large proportion of all utterances.
(2) The related words in a sentence often are not contiguous. The sentence 'the man who lived in the house sneezed' derives its meaning from the noncontiguous elements 'the man ... sneezed' and not from the contiguous elements 'the house sneezed'.
(3) Superficially different sentences such as the active `the boy hit the dog' and the passive `the dog was hit by the boy' have the same meaning.
(4) The same sound can have more than one meaning.*

It is hard to see how reference to any response, or partial response, or hidden partial response will remove the ambiguity associated with the use of sounds with two different meanings. Disambiguation almost always depends on the context within which the sound occurs. It is to explain the existence and importance of these features of language that Chomsky

* It is not clear whether one should speak of one word with two meanings or two words which sound alike (homophones).

and most subsequent writers on linguistics have postulated the presence of a deep structure in addition to the more superficial one which generates the left-to-right temporally ordered pattern of the observed behaviour.

The form of the deep structure. The deep structure was designed to account for the creative aspects of language, the connectedness between non-contiguous surface elements, the different meanings of a sound depending on its context, and the relationship between such superficially different sentences as the active and passive forms. In Chomsky's (1957, 1965) systems the deep structure consists of a complex set of rules which operate on symbols or strings of symbols, as well as the structures which the application of the rules generates. Chomsky's grammar has three sets of rules: the most important of these are *syntactic,* the others are *semantic* and *phonological.* The syntactic rules were seen as the creative part of the system, generating the basic sentence structure; the other rules acted passively on the inputs they received from the syntactic component to generate meaning and phonological representations.

Let us briefly consider the rules of the syntactic component of Chomsky's grammar and the structures that they generate. There are two different types of syntactic rule, *phrase-structure* rules and *transformational* rules. The phrase-structure rules operate upon symbols for grammatical categories such as noun phrase (NP) and verb phrase (VP), rewriting them as strings of symbols. Fig. 33 gives some examples of the operation of these rules. Note that these rules operate on individual symbols, and that no account is taken of the history of that symbol or of the derivation of the string in which it is embedded. A noun phrase receives the same treatment wherever it appears. The structure generated by the operation of the phrase structure rules is often portrayed as a tree diagram called a *phrase marker* (see Fig. 34); this seems to be viewed by Chomsky as a static structure, all parts of which exist simultaneously in the base component of the grammar.

Transformational rules were introduced when it was seen that although phrase-structure rules could generate simple active sentences, by themselves they could not account for such things as passive sentences or questions. The transformational rules are applied to the deep structures generated by the base component and differ from the phrase-structure rules in that they apply to strings of symbols, are applicable only in a fixed order, and take into account the history or derivation of the string upon which they operate. This last property means that, in a sense, they operate upon whole phrase markers. By allowing for the optional (Chomsky 1957) or obligatory (Chomsky 1965) addition, deletion, or re-ordering of elements within a string, they easily provide for the transformation of sentences from, for instance, active to passive or declarative to interrogative (see Fig. 35).

Rewrite SENTENCE as NOUN PHRASE (NP)+VERB PHRASE (VP)
S ⟶ NP + VP

Rewrite NOUN PHRASE as DETERMINER (DET)+NOUN (N)
NP ⟶ DET + N

Rewrite VERB PHRASE as VERB (V)+NOUN PHRASE (NP)
VP ⟶ V + NP

Rewrite NOUN (N) as dog, boy
N dog, boy

Rewrite DETERMINER (DET) as the
DET ⟶ the

FIG. 33. Examples of phrase-structure rules for the sentence `The dog bites the boy'.

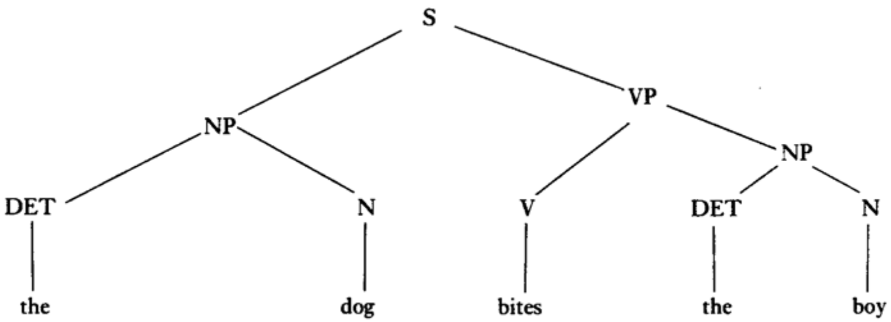

FIG. 34. An example of a phrase marker for the same sentence as in Fig. 33.

$\mathrm{DET}_{\mathrm{I}} + \mathrm{N}_{\mathrm{I}} + \mathrm{V} + \mathrm{DET}_2 + \mathrm{N}_2 \longrightarrow$

The dog bites the boy

$\mathrm{DET}_2 + \mathrm{N}_2 + \mathrm{BE} + \mathrm{V} + \mathrm{EN} + \mathrm{BY} + \mathrm{DET}_{\mathrm{I}} + \mathrm{N}_{\mathrm{I}}$
The boy is bitten by the dog

FIG. 35. An example of the transformation rule for converting active to passive sentences. Same sentence as in Fig. 33.

In Chomsky's theory the meaning of a sentence is assigned to it by a separate component, the semantic component. This acts passively on the terminal string of the phrase marker, fitting meanings to each of the elements. Meanings are fully determined by the nature of the input from the syntactic component to the semantic component. Katz and his colleagues (Katz and Fodor 1963, Katz and Postal 1964, Katz 1972) have constructed a theory for this kind of semantic component, envisaging it as composed of two parts: a dictionary of meanings, and a set of projective rules allowing for, and providing meanings of, combinations of items.*

Chomsky's system succeeds in doing what it set out to do. It accounts for many of the interesting features of languages which fall outside the province of simple behaviourist models; it permits the generation of an infinite number of sentences from a finite number of rules; it explains why distant elements of a sentence can have strong relationships; it answers to our intuitive feeling of a similarity between syntactically different sentences by identifying a common deep structure. However, as Chomsky himself noted (1965, p. 162), it fails in one important respect; it does not capture the still deeper semantic relationships which can exist between syntactically different sentences. Thus, the grammar fails to capture the similarity between sentences A and B, or C and D:

(A) I liked the play.
(B) The play pleased me.
(C) John bought the book from Bill.
(D) Bill sold the book to John.

This failure to account for *paraphrases* would appear to be due to the narrow definition of the semantic component of the system. As we have seen, it is purely a passive feature of the grammar, whose function is to ascribe meaning to the deep structures generated by the base component. Intuitively, this seems to be an unnecessary restriction on the role of the semantic component. The meaning of a sentence is *not* only the sum total of the meanings of the words but includes the way in which they are put together. In this broader sense of semantic the base structure itself should be included in the semantic, and not in the syntactic, component.

This broader usage of semantic requires some elaboration, for it embodies an important shift in thinking about language comprehension. The behaviourist emphasis upon the elements of speech meant that most research was concerned with individual items; how they were processed,

* The dictionary operates on the basis of *componential analysis*, specifying the meaning of a lexical element as the set of categories within which it is included (*semantic markers*) together with those features which separate it from other lexical items in the same categories (*distinguishers*). The technique of componential analysis has also been applied to verbs, in particular by Bendix (1966, 1971). He examined a number of verbs and showed that they could all be paraphrased by combinations of a few basic verbs such as have, cause, change, etc.

stored, interpreted, and generated. Thus, standard experiments involved the learning of lists of words or paired associates. When organizational factors were allowed (cf. Mandler 1967), they were generally restricted to the meaning relationships between isolated words. This accounts for the notion of *categories* and the host of experiments on the role of categorial relations in the learning of word lists and paired associates. Chomsky's critique partly embodied the notion that lexical items were not the central elements in language comprehension. However, in moving to higher-order units Chomsky did not expand the semantic component to include the meaning of these larger units.

Most of the recent work concerned with semantic deep structures, then, concentrates upon the mechanisms for comprehending and storing these higher-order verbal units, beginning with the recognition that what is remembered of sentences, paragraphs, or even stories is the sense of the discourse as a whole. Before turning to a brief discussion of some of this work it is worth digressing momentarily to discuss Tulving's (1972) notion of semantic memory. In view of the shift we have just described, it is unfortunate that Tulving chose to apply the term semantic in its older usage to a system representing the meaning of individual words, independent of context. The confusion arising from this usage has led some (e.g. Schank 1975) to reject the notion of semantic memory entirely in favour of a system including only the lexicon and episodic memory. According to Schank, the meaning of individual words is stored in the lexicon, while any relations between individual items must be stored in terms of some event in which they took part. We cannot agree with Schank on this point, though we find his model for semantic deep structure (see p. 398) one of the most attractive in the field. While we do not accept Tulving's separation *in toto,* we think there is strong evidence for a separation between some form of context-free memory, using (in the old sense) semantic categories, and a context-dependent memory, using something like a spatio-temporal framework. As we shall see in the next chapter, the data from amnesic patients supports this distinction.

Semantic deep structure. The work of Bransford, Franks, and their colleagues (e.g. Bransford and Franks 1971, Bransford, Barclay, and Franks 1972, Bransford and Johnson 1973, Bransford and McCarrell 1974, Franks 1974) provides important clues to the nature of memory for higher-order verbal units. Their early work demonstrated that, given a set of related sentences, subjects formed something like a prototype sentence which, though never actually seen, was more readily recognized as familiar than sentences which had been seen. Later work extended this observation by showing that the remembered representation for a sentence depended upon the context within which it was seen, as well as upon various inferences and assumptions the subjects could make about the material,

presumably based on some prior knowledge of the contexts within which the events described could obtain. In fact, given an inappropriate context, a sentence which would have been understood in isolation was often judged incomprehensible. Similarly, they argued, some sentences which would be meaningless in isolation can be given some sense by the context within which they occur.

This work on sentence comprehension requires a model which provides for some deep structure that codes the relationship between the various elements in the sentence or between several sentences. Studies of semantic deep structure have concentrated upon such models, in the hope of specifying the form within which these relationships could be coded such that the meaning of a sentence as a whole could be stored, paraphrases of that sentence recognized, sentences could influence one another's representations, and prior information could be brought to bear on comprehension of inputs (and, hence, the meaning attached to these).

Early work on the basis for a semantic deep structure (e.g. Bendix 1966, 1971, Fillmore 1968, 1971, McCawley 1968, 1971, Lakoff 1971) spoke primarily to the first two of these requirements, concentrating upon sentence comprehension in isolation. Though superseded by later models, we shall describe Fillmore's system as it presents some of the basic features of those which superseded it. According to Fillmore, a *case system,* in which items were unordered though identified as to function, would provide a more appropriate base than the ordered set of grammatical categories proposed by Chomsky. In Fillmore's system the sentence is represented by its *modality,* which specifies such conditions as tense, negation, and mood of the sentence as a whole, and *proposition*, which identifies the verb and its permissible cases. These latter are given as an unordered set, with each case defining the relationship between the item in that case and the verb. Fillmore specified eight deep-structure cases:

(*a*) Agent—the instigator of the event
(*b*) Counter-agent—the resistance against which the action occurs
(*c*) Object—the entity acted upon or under consideration
(*d*) Result—the entity that ensues from the action
(*e*) Instrument—the immediate cause of the action
(*f*) Source—the place from which some entity moves
(*g*) Goal—the place to which some entity moves
(*h*) Experiencer—the entity receiving, accepting, undergoing, or experiencing the effect of an action

In the sentence *John opened the door with the key,* John is the agent, door the object, and key the instrument. The deep structure of each simple sentence would consist in a verb plus its obligatory and optional cases. *Open,* for example, always requires an object, but takes an agent and an instrument as options. The transformational rules in Fillmore's system, as in Chomsky's, are concerned with generating surface sentences from deep structures. However, since the cases in Fillmore's semantic deep structure are unordered, there is no need for rules which transpose elements. Instead, the rules establish a hierarchy amongst the cases associated with a verb, specifying which grammatical role each case will play in the surface sentence. For *open,* the instrument is the subject if it occurs alone, but the object of a prepositional phrase (with the *key)* if there is an agent. Fillmore's grammar does require deletion rules, because cases are represented in the deep structure by prepositional phrases which, in most circumstances, do not appear in the surface structure. Thus, the agent in our example would be represented in deep structure as *by John.* The preposition would survive in the surface sentence only in the passive case; in the active form the *by* would be deleted by a transformation rule.

This type of semantic deep structure, important for its emphasis upon functions and actions, can account for many of the facts of sentence comprehension. However, it remains silent on the more complex problems delimited by Bransford, Franks, and others, and those represented by the retention of the sense of paragraphs or entire stories. Three recent models which are specifically pitched at this level seem particularly interesting, those of Schank (1972, 1975), Norman and Rumelhart (1975) and Jackendoff (1976). Common to these approaches is the assumption that the deep-structure representation for language is some form of propositional or conceptual network which codes meaning through the interaction of elements. Thus, for Schank (1975) the basis of human memory is the *conceptualization,* which is 'action-based with certain specified associative links between actions and objects' (p. 259). Similarly, for Norman and Rumelhart (1975) the basis is the *active structural network,* which is a semantic network representing the underlying propositions in any stored event. Both systems rely on a set of primitives which define the forms of interaction between the elements in the memory structures; here they follow in the path of Bendix's componential analysis of verbs. Further, both argue that sentence after sentence can be 'added' to the memory structure, in some cases being influenced by what is already there, in other cases influencing it. Thus, they provide models for the comprehension of sets of sentences. More recent work by Rumelhart (1975) attempts to provide the basis for a representational network which would describe the structure of an entire story without building sentence upon sentence.

While we cannot explore these models in detail, it is worth emphasizing the fact that they insist upon a network-like propositional representation where the elements within the network are related to one another through the action of a primitive set of operators. The meaning of such a network, or conceptualization, is the totality of the relationships embodied within it. We find it particularly heartening that Norman and Rumelhart

emphasize the essential non-linguistic character of their networks; they apply their analysis to imagery phenomena as well as to linguistic deep structure. Here, they also stress the view that imagery depends, not on a pictorialization within memory, but rather upon some propositional deep structure which captures the relationships embodied in the image and from which the image can be reconstructed.

We will conclude this section on deep structure models with a discussion of Jackendoff's system (Jackendoff, 1976) which is, for us, the most interesting and exciting of the recent proposals. Jackendoff, expanding an original suggestion by Gruber 1965, has proposed that all sentences have deep semantic structures which are formally analagous to the subset of sentences describing events or states of affairs in physical space. First he shows how an analysis similar to the one by Fillmore described above will provide a deep structure for sentences about the location and movement of entities in physical space and, second, he shows how modifications and extensions of this purely spatial system can account for the meanings of non-spatial sentences.

In his analysis of spatial sentences, he starts with examples like:

(1) The train travelled from Detroit to Cincinnati
(2) The hawk flew from its nest to the ground
(3) The rock fell from the roof to the ground

and shows how their meanings can be captured by a deep structure which specifies the thematic relations between the verb and the nouns or noun phrases. Thus (1) would be represented by the deep structure function, GO; the theme of the function, *train;* the source or place from which the movement started, *Detroit;* and the goal or place where the movement ends, *Cincinnati.* Notice the similarity to Fillmore's case system described above. Spatial sentences (2) and (3) above would have similar deep structures with suitable additional information such as the manner of the motion. Other spatial sentences such as

(4) Max is in Africa
(5) The cat lay on the couch
(6) The bacteria stayed in his body
(7) Bill kept the book on the shelf

describe not the motion of the object or theme but its location and are represented by the deep structure function BE (4 and 5) or STAY (6 and 7).

Thus all states of affairs and events in physical space can be represented in Jackendoff's system by three functions GO, BE, and STAY, together with the things and places which these functions relate. Agency and causation are added to the deep structure by the higher order functions CAUSE and LET which apply not to entities but to events. Thus if sentence (3) above would be represented by GO (THE ROCK, THE ROOF, THE GROUND), then

(8) Linda threw the rock from the roof to the ground
(9) Linda dropped the rock from the roof to the ground

would be represented as CAUSE (LINDA, GO (THE ROCK, THE ROOF, THE GROUND)) and LET (LINDA, GO (THE ROCK, THE ROOF, THE GROUND)) respectively.

At this point Jackendoff takes a crucial step. He claims that nonspatial sentences have exactly the same representation except that the functions GO, BE, and STAY do not refer to the spatial location of entities but to the possessive, identificational, or circumstantial 'location' of entities. Let us look at possessive GO. While spatial or positional GO signifies the movement of an entity from one physical location to another, possessive GO signifies the movement of an entity from one possessive location to another. The sentence

(10) Harry gave the book to the library

is represented as possessive GO (THE BOOK, HARRY, THE LIBRARY). Similarly

(11) The book belonged to the library
(12) The library kept the book

are represented by possessive BE (THE BOOK, THE LIBRARY) and possessive STAY (THE BOOK, THE LIBRARY). The analysis of sentences about continuing states of identity or changes of identity, or continuing or changing circumstances, are given a similar treatment. Jackendoff shows, for example, that the semantic analysis of the sentence

(13) Linda kept Laura from screaming

is exactly parallel to the sentence about physical prevention

(14) Linda kept Laura (away) from the cookie jar

except that the avoided location is a circumstance in example (13). Thus, for example, the same rules of inference allow us to conclude that Laura did not scream (13) nor did she get to the cookie jar (14).

Jackendoff has not extended his analysis into the domain of completely abstract concepts nor to verbs referring to internal states or beliefs, but he sees no insurmountable obstacle to such a programme, nor do we. It is also reasonable to assume that this type of analysis can be extended to deal with units of speech longer than a sentence, thus incorporating the recent work on discourse and narrative comprehension. In summary, Jackendoff says:

'I consider it a striking property of the present system that simple principles, framed in terms of physical space, can be stated formally in such a way as to generalize to domains that bear no *a priori* relation to physical space' (Jackendoff, 1976 p. 121).

On linguistic grounds, then, it appears necessary to postulate the existence of several different mechanisms underlying the production and understanding of language. In addition to those mechanisms which select the appropriate words and sentence frames to produce the temporal left-to-right ordered structure of the surface aspects of language, there must be a deeper, more abstract level which carries the sense of a sentence or a set of sentences.* The common element that all deep structures share is their non-temporal aspect; put another way, they can all be represented by purely static spatial structures. The sense of an item is derived from its relation to other items within the structure, the overall sense of the sentence follows from the total configuration, while the interaction between such configurations allows for higher-order messages such as stories.

In terms of our model all of these deep-structure elements are identified with maps in the locale system (or their activation). The surface structure of the grammar, transformational processes, the syntactic structures, and the lexicon are those parts of the taxon system which provide the means by which maps in one person's locale system are transferred to another's.** These taxon systems are analogues of the route systems of lower animals. They are based on the operations of categorization and the formation of linkages between frequently associated items to yield route statements. These routes, which take discourse from one substantive to another, would appear to be the basis for the tone groups described by Laver (1970, see footnote below). The rules which govern the generation of a set of routes from the underlying semantic map are analogous to the transformational component of a Chomsky-type grammar. As we shall see these rules can be much simpler than those specified by Chomsky or Fillmore, since the form of the surface sentence can be read directly from the directions traversed in the map.

Types of semantic maps. In developing the notion of a semantic map we shall build on Jackendoff's insights and follow his general methods. First we shall describe how semantic maps can be used to provide semantic deep structures for sentences about entities and events in physical space. Using one of these spatial semantic maps as an illustration,

* Studies of language generation, which occurs in small segments called tone groups according to Boomer and Laver (1968) and Laver (1970), indicate that there are three separable processes involved (Goldman-Eisler 1968): (1) an idea or determining tendency; (2) the transformation of this idea into a sequential chain of symbols; (3) the selection of appropriate lexical items. One function for the verbal short-term memory system noted before (p. 388) could be the retention of a group of surface elements during the elaboration of the entire tone group. In contrast to the appearance of lexical elements in a short-term holding system, both at input and output, long-term memory for language is clearly concerned with the meaning of an utterance (e.g. Johnson-Laird 1970, Bransford and McCarrell (1974).

** From an evolutionary point of view, language could have developed as a means of transferring information about the spatial aspects of the environment: how to get somewhere, how to find food, etc.

we shall outline some of the syntactic transformation rules for transcribing all or part of this map into a surface 'route' sentence or phrase. Following Jackendoff's method, the next step will involve a discussion of how non-spatial maps similar to these spatial maps can be formed and how the same transformation rules operate to generate sentences about non-spatial entities and events. Instead of mapping physical space, these non-spatial maps depict surfaces on which the locations represent possession (or, as we shall call it, influence), identity, and circumstances. In this section we will also introduce the notion that maps or parts of maps can be 'named' and these names can be entered into locations on other maps.

A semantic map for spatial sentences. Sentences about the location of objects or the occurrence of events in physical space have an obvious and natural representation within a spatial map structure. Let us use Jackendoff's sentence (3) as an example

(3) The rock fell from the roof to the ground

The map of this event has three phases (see Fig. 36). The first (a) depicts the unstated presupposition that an entity (E) *the rock* was in a place (P) *the*

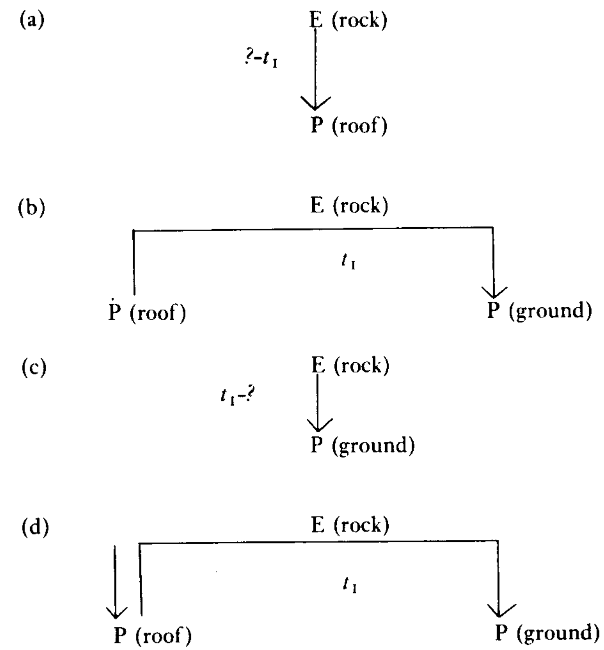

FIG. 36. Schematic for a semantic map of the physical spatial sentence 'The rock fell from the roof to the ground'.

roof up to some unknown past time t_1. The second (b) depicts the action at time t_1 in which an entity *the rock* moves from the place *the roof* to the place *the ground.* In the third phase (c), *the rock* is in the location *the ground* from time t_1 onwards. Thus there are two places, an entity which either stays in these places or moves from one to the other, and time markers which specify the time of the movement and the beginning and end of the period spent in a location. These time markers may refer to times attributed to the external world or they may be entirely internal to the map. The entire map is shown in Fig. 36(d).

Notice how the mapping system incorporates the three fundamental functions of Jackendoff's system, GO, BE, and STAY. BE is represented by the location of an entity in a place without a time marker. If there were no time marker on the first phase (a) of our three phase representation, this would depict the BE function: the rock is on the roof and, as far as we know, always was and always will be. The STAY function is represented by the third phase (c) where the time marker *t,* limits the duration of the state in the past direction but there is no time marker to limit it in the future direction. The second phase of our semantic map (b) represents Jackendoff's GO function, the movement of an entity from one location to another at some specific time.

A variety of sentences can be generated from our simple spatial semantic map by a set of transformation rules. We will assume that there are no obligatory points of entry into a semantic map nor are there obligatory directions of movement within the map. Thus although the map may have been originally constructed from a simple active sentence, it can be entered at any entity, event, place, or movement and read in part or in whole in any direction. The order of reading and the relationships between the successive items read determines the syntactic role of each item in the surface sentence. Maps containing nothing but spatial entities and events can only generate sentences in the active voice. We shall discuss the passive voice shortly. If the map of our example (Fig. 36) is entered at the entity *rock* and this is read first, it becomes the syntactic subject of an active sentence. If the movement is read next it becomes the verb, and the place of origin and the place of termination of the movement are made the objects of the prepositions *from* and *to* respectively. With this order of selection, we have generated our active sentence from which the map was created.

(15) The rock fell from the roof to the ground

If after reading the entity *the rock,* we had read first the place of origin *the roof* and then the movement and place of termination, our sentence would read:

(16) The rock was on the roof and (then) it fell to the ground (from there)

Similarly:

(17) The rock was on the ground where it had fallen from the roof

The natural expression of the relationship between an entity and its location is

$$\text{(Entity) is } \left.\begin{matrix}\text{in}\\ \text{on}\end{matrix}\right\} \text{ (location)}$$

We might have entered the map at one of the places and read the entity next.

(18) The roof had a rock (on it) and (then) the rock fell to the ground
(19) The ground has a rock (on it) which fell from the roof

The natural expression of the relationship between a location and its entity is

(Location) *have* (entity)

Similarly:

(20) The roof had something fall from it on to the ground and that was the rock
(21) From the roof, the rock fell to the ground

Finally the map can be entered at the movement itself

(22) The falling of the rock from the roof to the ground

In which case the movement is nominalized as a gerund or a noun (fall) and the entity becomes the object of the preposition *of* as a subjective genitive. It is clear, therefore, that any of the parts of a map can be the subject of the sentence derived from that map and that the parts of the map can be read in any order or direction. Finally the syntactic function of a representation in the map is determined partly by its role in the map and partly by the order in which the parts of the map are read.

The simple physical semantic map can embody all three of the deep structure functions (GO, STAY, BE) which were discussed in the previous section. Note that there are no rules necessary to derive inferences from the deep structure since these are built into the map when it is constructed in the first place. Our map of sentence (15) contains the unstated information that the rock was on the roof for some undetermined time, that it fell through the places between the roof and ground before falling to the ground, and that it remained on the ground for some undetermined time after the event.

Semantic maps for non-spatial sentences. Sentences about entities and events in physical space constitute only a small proportion of the language. Non-spatial sentences represent notions such as possession, causation,

responsibility, identity, and category inclusion to name a few. Jackendoff introduced the higher order functions CAUSE and LET into his linguistic system to deal with causation and permission in spatial sentences. More importantly he suggested that non-spatial sentences conveying information about possession, identity, and circumstances had the same formal structure as spatial sentences. We will draw upon this insight and introduce the notion of non-spatial maps. Following Jackendoff, we will consider three types of non-spatial map: maps of influence or possession surfaces, maps of identity surfaces, and maps of circumstance surfaces. We do not have room here to go into great detail about each of these non-spatial maps so we will concentrate on the influence or possession type and only briefly comment on the other two.

One important concept not captured by a purely spatial mapping system is that which is common to the notions of causation, control, power, instrumentality, and possession. These notions represent a relationship between entities and/or events in which one is under the influence of the other. Some of these relationships are represented in semantic systems such as Fillmore's by the deep semantic cases *agent* and *instrument* (see above). We will postulate that all of these relations are represented on one surface which we shall call an influence surface. Influence relations on this surface are represented by entities in particular locations and changes in influence are portrayed as movements between places. Expanding on our previous notation, places in our influence map will be labelled P_{infl} while places in our physical spatial map will be labelled P_{phys}.

Entries in different maps (entities, places, and movements) which have the same name are considered to be connected so that the activation of one entry also activates all the other entities. For example, if *Harry* were the name of a place in influence space and an entity in physical space, activation of one would activate the other.

Before we discuss an example of a map portraying relations and events in an influence space, it will be useful to introduce the concept of *map nesting* or *embedding.* Maps or parts of maps can be labelled with names and these names can then be represented as entities or locations in other maps. The names of maps can appear not only in maps of the same type but of other types as well. Thus the name of an influence map could appear as an entity in a physical spatial map. This notion of map embedding will become clear in the next example which illustrates both the movement of an entity in an influence space and the embedding of this influence event in a second influence map. Our example is sentence (10) taken from Jackendoff (see above).

(23) Harry gave the book to the library

This means that the book moved into the possession of the library and that this event was caused by Harry. In the system we are proposing both the transfer of possession and the causation of that event would be represented in an influence map. The transfer of possession is represented as a movement from some unknown location into the location (the library). This event is given the name *transfer of possession* and entered into the location (Harry) in the influence map. The interpretation given to the relationship in an influence map between a location and its content depends on the nature of the content. When the content is a primitive entity drawn from a taxon store such as *the rock* or *the book* then the relationship is interpreted as a possession where the location possesses the content. When the content is the name of another map (i.e. an event), then the relationship is one of causation, the event is caused by the location. Entities possess other entities, entities or events cause events.* We will leave open for the moment what the interpretation of an entity in an influence location of an event might be.

Fig. 37 shows the influence map of sentence (23). Notice that the sentence does not specify whether the book belonged to Harry before the event described or whether the book actually physically moved to the library. These are left ambiguous. Consider the related sentences:

(24) Harry gave his own book to the library

which disambiguates the book's former possessor.

(25) Harry gave the book to the library- after it had been displayed on loan there for several years.

and

(26) Harry gave the book to the library but won't be sending it to them until next year.

The map representation of the event in sentence (23) is given in Fig. 37 (a)-(d). Notice the similarity to Fig. 36 except that the relationships and the event take place in an influence space. In order to represent the rest of the sentence, namely, that Harry caused the event portrayed in Fig. 37(d), this influence map is given a name and entered into the influence location called Harry (Fig. 37(e) and (f)). Fig. 37(g) shows the map of the whole sentence. As we stated earlier, entities *are possessed* by influence locations (e.g. Fig. 37(c)) but events *are caused* by influence locations (e.g. Fig. 37(f)).

More than one entity can be represented in an influence location. For example, if our influence map had represented sentence (24) instead of

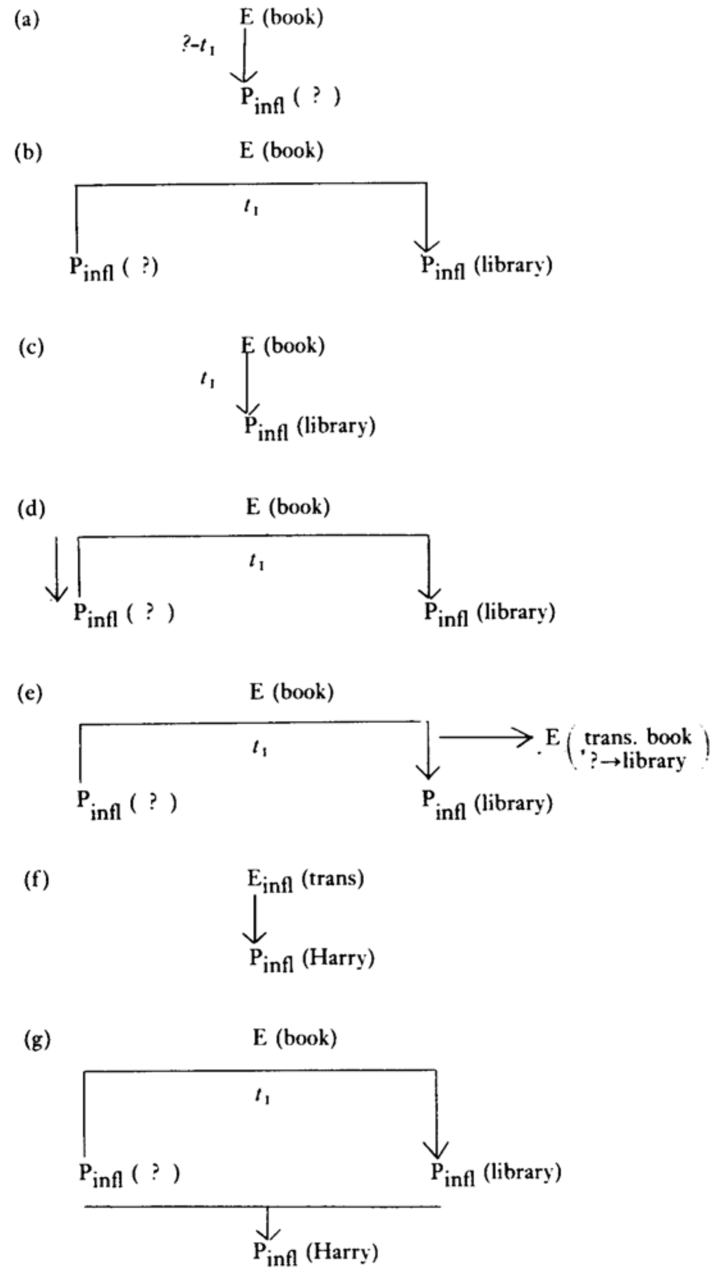

FIG. 37. Schematic for an influence map of the sentence 'Harry gave the book to the library'.

* Strictly speaking, only events *cause* other events. When an event is entered into the influence space of an entity this is interpreted as an instrumental relationship. We will not go into this complication but assume for the present that agentive entities can cause events. It does not change the basic arguments set out here.

sentence (23) then the location *Harry* would have been substituted for the unknown location of the book in Fig. 37(a) and this location would have contained both the book until time t_1 , and the transfer at time t_1.

Let us look at some of the transformation rules for reading route sentences from influence maps. Here we will continually refer to the rules developed for physical spatial maps to show the similarities. As with physical spatial maps, influence maps can be entered at any point and read in any direction. Let us start with the event in influence space and compare it with the event in physical space mapped in Fig. 36. Consider the following pairs of sentences:

(15) The rock fell from the roof to the ground
(27) The book went from (the possession of) someone to (the possession of) the library
(16) The rock was on the roof and (then) it fell to the ground
(28) The book was in the possession of someone (or was someone's) and (then) it went to the library
(17) The rock was on the ground where it had fallen from the roof
(29) The book was in the possession of the library whence it had come from some unknown person

The natural syntactic expression of the relationship between an entity and its location in influence space is the same as in physical space. Conversely the natural expression of the relationship between an influence location and its content is

(Location) *have* (entity)

(18) The roof had a rock (on it) and (then) the rock fell to the ground
(30) Someone had a book and (then) the book went to the library
(31) Harry had the book given to the library
(32) Harry had someone give his book to the library

As examples (31) and (32) show the *have* transformation holds irrespective of whether the content of the influence location is an entity possessed or an event caused.

Two important syntactic features which must be introduced into the transformation rules for influence maps do not exist for physical spatial maps. These are the active/passive voice distinction and the genitive of possession. Both of these are necessary to transcribe readings in which the contents of an influence location are read before the name of the location itself. Thus, in our example, reading the locations first in Fig. 37(c) and (g) gives, respectively

(33) The library's book
(34) Harry had the book given to the library

while reading the contents of those locations first gives

(35) The book of the library
(36) The book was given to the library by Harry

If the influence location is read after its contents, it becomes the object of the preposition *by* in a passive sentence or the object of the preposition *of*. Notice that instead of (34) one could have read (37) or (38):

(37) Harry's giving (of) the book to the library
(38) Harry's gift of the book to the library

In addition to a mode expressing possessional relations and changes in possessional relations, Jackendoff postulated *identificational and circumstantial* modes whose syntax was analogous to the physical spatial mode in much the same way that that of the possessional mode was. We have not attempted to work this out in detail or to construct maps of these spaces* but see no insurmountable problem in doing so. Note, for example, the similarities between the identificational sentence (39), the circumstantial sentence (40), and the physical sentence (15) above.

(39) The rock went from smooth to pitted
(40) The librarian went from laughing to crying

To conclude this section, let us briefly mention what seems to us the main weakness of the semantic map idea. Maps of physical space which use a Euclidean metric allow inferences to be drawn about the relationships amongst entities in those maps which go beyond the usual laws of logic. This was, of course, what Kant meant when he called space a *synthetic a priori.* Thus we can say that the rock in its fall from the roof to the ground must pass through the intervening places. Our map did not specify whether these places were occupied or not. If it had then we could conclude that the rock must bump into, or pass through, whatever occupied those places. Some of the work of Bransford and Franks and their colleagues seems to be based on inferences of this type. Now it is not immediately obvious that the same kinds of inferences can be drawn from our non-physical semantic maps. Part of our difficulty here is that we do not know what the axes of the non-physical semantic maps are or even whether the maps can be considered Euclidean. These objections notwithstanding we think that the work of Jackendoff opens up exciting possibilities for investigating the use of maps of physical and non-physical spaces as the basis for deep semantic structures.

14.3.3. CONCLUSION

The above sections outline a possible way in which the cognitive-mapping system could function as a deep structure for language. The representations in this deep structure have all the properties we have attributed to maps.

* Whether these maps are irreducible primitives or will on further examination be found to be reducible to various combinations of physical spatial and influence maps remains to be seen.

Items are entered into a semantic map not on the basis of the order in which they are received or their position in a left-to-right linear string but in accordance with their semantic relationships to other items in the map. Maps do not incorporate information about the way that they were constructed or the 'route' surface sentence from which they were constructed. Items in the map and their relationship can be read in any order and a large number of different sentences can be generated from the same map. Once a semantic map exists, additional information can be added or changes in the existing information can be made. The resulting large semantic structures embody the ideas expressed in paragraphs or whole stories. We locate these semantic maps in the left human hippocampus. Taken together with the representations of physical occurrences in the right hippocampus they form the basis for what is generally referred to as long-term, context-specific memory for episodes and narratives. In the next chapter, this assertion will be 'tested' against the known facts of the amnesic syndrome.