(1988). In J.W. Berry, S.H. Irvine, and E.B. Hunt (Eds.) Indigenous cognition: Functioning in cultural context, (Pp. 57-84). Boston: Martinus Nijhoff Publishers.

THE WHORFIAN HYPOTHESIS REVISITED: A COGNITIVE SCIENCE VIEW OF LINGUISTIC AND CULTURAL EFFECTS ON THOUGHT

Earl Hunt and Mahzarin R. Banaji

Therefore is the name of it called Babel; because there the Lord did confound the languages of all the earth..... Genesis; 11:9.

When the people of the earth ceased to have the same language, they lost the ability to communicate. But did they continue to have the same thoughts, expressed in different tongues? We think not. Consider a more modern failure to communicate. The historian Barbara Tuchman has admitted that she simply cannot write about certain types of people.

> Not ... a cleric or saint, for they are outside the limits of my comprehension. (Tuchman, 1978, p. xiv)

Why are fourteenth century clerics outside of the comprehension of an extremely erudite twentieth century woman?

We believe that virtually everyone is agreed that culture does influence thought. There is also a widely held intuition that language is important. Benjamin Lee Whorf (1956) presented this argument so elegantly that the intuition is often referred to as the "Whorfian Hypothesis." Whorf argued from his own observations and well chosen examples. Controlled observations, however, have generally failed to give very much support to what seems to be a reasonable idea. Why? In this paper we shall re-examine the logic of the Whorfian hypothesis, from the viewpoint of modern cognitive psychology. More specifically, we shall maintain that modern theories of cognition imply the Whorfian hypothesis, in a modified form, and restrict its influence in an orderly way. Thus, we go beyond Whorf in presenting a model of how language acts on thought, and by using the model, to state limits on the influence of language.

Our argument will be presented in three stages. The section immediately following presents a summary of the Whorfian hypothesis and related theoretical and empirical work. The next section describes what we

believe to be a reasonable model of mental information processing, given the current state of cognitive science. The third and fourth sections unite the two, by presenting examples of how thoughts are produced by the interaction between linguistic knowledge and information processing mechanics. We will follow Whorf's tradition by arguing from example, rather than by following the experimental psychologists' tradition of controlled observation. The final section of this paper is a summary and commentary.

The Whorfian Hypothesis

The concept of linguistic relativity is central to Whorf's hypothesis. This concept had been proposed by Whorf's mentor, Sapir (1941), who took the strong position that language imposed perception upon reality. In his own words,

> The fact of the matter is that the "real world" is to a large extent unconsciously built up on the language habits of the group ... We see and hear and otherwise experience very largely as we do because the language habits of our community predispose certain choices of interpretation. (Sapir, 1941; also in Whorf, 1956, p. 134)

Although Sapir's ideas attracted attention, he was unclear about the nature of the evidence required to confirm his hypothesis. Whorf published two papers, *Science and Linguistics* and *Linguistics as an Exact Science*, that attempted to fill this gap. In these papers, he claimed that all higher order thinking is dependent on language. Whorf's restatement of linguistic relativity was,

> We are thus introduced to a new principle of relativity, which holds that all observations are not led by the same physical evidence to the same picture of the universe, unless their linguistic backgrounds are similar, or can in some way be calibrated.

and that,

> users of markedly different grammars are pointed by their grammars toward different types of observations and different evaluations of externally similar acts of observation, and hence are not equivalent as observers but must arrive at somewhat different views of the world. (Whorf, 1956, pp. 214 and 221)

To prove his case, he offered numerous examples contrasting "Standard Average European" (SAE) thinking to thinking in the Hopi and Shawnee languages, which he had studied on field trips. He also offered numerous examples from his own professional experiences. Whorf had worked as an insurance inspector for fire safety standards. He noticed that workers would smoke near drums filled with fumes more often than those filled with gasoline, even though the former were more dangerous. Whorf's analysis was that ,

> Physically the situation is hazardous, but the linguistic analysis according to regular analogy must employ the word "empty", which inevitably suggests lack of hazard. (Whorf, 1956, p. 134)

Another example further develops the idea that behaviour is influenced by the constraints of the linguistic formula. While examining a wood distillery, Whorf noted that no precaution was taken to cover the limestone used for insulation from contact with flame, even though flammable acetic acid deposits were building up on it. Distillery workers were surprised when the "limestone" began to burn. Again, the label "limestone" had been misleading, because "stone" implied noncombustability. We shall offer a more detailed discussion of such examples in the following section.

An impressive paper contained in a collection of posthumously published works (Whorf, 1956), *"The relation of habitual thought and behaviour to language"* addresses the question: "Are our concepts of 'time', 'space', and 'matter' given in substantially the same form by experience to all men, or are they in part conditioned by the structure of particular languages?" To answer, Whorf turned to the contrast between European and Hopi linguistic treatments of time, space, number, and sequence. Here are two of his examples:

(1) In English there are two types of nouns to denote physical objects: the individual nouns (for example, a chair, a clock, a computer, and a book) and mass nouns (such as water, soup, sand, and flour). In Hopi, there is no formal subclass of mass nouns. Instead, the noun for different forms of the object implies the specific form. English speakers would define a form for water by defining a container, as in "a glass of water" or "a pool of water". The Hopi would use a different word for each form.

(2) The Hopi have a large vocabulary of terms to express duration and intensity. This is because they do not make use of physical metaphors. Whorf observed that English speakers may say,

> I "grasp" the "thread" of another's argument, but if its "level" is "over my head" my attention may "wander" and "lose touch" with

> the "drift" of it, so that when he "comes" to his "point", we differ "widely", our "views" being indeed so "far apart" that the "things" he says "appear" much too arbitrary, or even "a lot" of nonsense! (Whorf, 1956, p. 141)

The Hopi could not use verbs metaphorically, because in Hopi, verbs describing physical actions can only appear in their literal context. In order to express a thought like that offered above, the Hopi would use a special class of "tensor" words to express intensity, duration and tendencies of thought. As a result, the Hopi would stress the development and decline of an event. This was reflected in the cultural importance of ceremonies such as meditation to prepare oneself for an event and announcement that an event had progressed to a new stage.

How do these differences in grammar between the Hopi and SAE translate into differences in thought processes? We shall answer this question by offering our own interpretation of Whorf's ideas. He believed that speakers of European languages analyse the world in terms of things that have a unique location in space. To further structure the world into discrete categories, nonspatial events are given attributes of form and continuity. For the Hopi, the world is analysed in terms of events whose different parts are strongly interactive if they occur at the same time. We will illustrate by taking one of Whorf's examples, a rosebush. From the Western point of view a rosebush is a thing, with its unique location, that is distinct from other things in different locations. In surprisingly modern terms, Whorf (1956, p. 150) points out that when Western people (cognitive psychologists?) think of a rosebush, they believe they are manipulating a mental image that represents a rosebush, but that is distinct from it. On the other hand, a rosebush is also a process that buds, flowers, and decays. The Hopi would see their thought as an event that was coterminous with and influencing the processes of change in the rosebush itself.

Whorf believed that these different modes of thought are, if not dictated by, at least strongly influenced by the differences between SAE and Hopi languages. As the Hopi do not have words to express a thing-like metaphor for the rosebush, they cannot think about it as a thing, it is a process. As we write this, we have difficulty expressing what the Hopi would have thought, because we must express their idea in the inadequate English language and, perhaps, because our own thought is constrained by English.

Note that we have said "constrained" and not "dictated." This is the crux of the controversy about Whorf's ideas. We believe that Whorf was a linguistic relativist, not a linguistic determinist. He did not believe that thought was dictated by language, but he did believe that language predisposed thoughts to take certain shapes. Consider his views of science:

> ... the world view of modern science arises by higher specialisation of the basic grammar of the Western Indo-European languages. Science of course was not caused by this grammar; it was simply coloured by it. (Whorf, 1956, pp. 221-222)

The problem with being a linguistic relativist is that the category name is not sufficiently constraining. What are the boundaries of language's influence on thought, and how are these boundaries established? Under what circumstances can a person override the boundaries of his or her own language to understand the concepts of a foreign culture? We shall attempt to answer these questions by presenting a general model of human thought, showing that the model implies a form of the Whorfian hypothesis, and by developing principled restrictions on the hypothesis itself.

A Model of Mental Mechanics

Our view of mental action is based upon a rather sharp distinction between two aspects of thought: thoughts as a process of internal symbol manipulation independent of the meaning of the symbols; and thought as the manipulation of an internal representation of a (real or imagined) external situation. The distinction has been presented in some detail elsewhere (Anderson, 1983; Hunt, 1983; Newell, 1980; Pylyshyn, 1984), so we shall deal with it only briefly. In common with most cognitive scientists, we regard "thinking" as a manipulation of an internal model of the world. As an abstract computation, this manipulation must follow species-general, culture-free laws. For instance, we assume that the process by which information is moved from short term memory to permanent memory is the same in everyone, although we would allow for some individual variation in the efficiency of the process. On the other hand, the content of the information acquired from a particular experience will be influenced by those aspects of the situation on which a person chooses to "fix attention", i.e., to bring into memory in the first place. Thus, the content of the information acquired will, in general, be culture-specific.

For brevity, these two aspects of thought will be referred to as the *mechanistic* aspect and the *representational* aspect of cognition. The mechanistic aspect is quite outside our conscious experience, although models of mechanistic thought can be evaluated by experimental observation. Otherwise experimental psychology would be impossible. The representational aspect is at least partially part of our conscious awareness.

To illustrate, if a person's actions remind you of a gorilla you are aware of thinking of the gorilla, but quite unaware of how you thought of it.

A complete model of mechanistic thought would be quite detailed. Models to account for only a few classes of experimental observation have been published by Anderson (1983), Hayes-Roth and Hayes-Roth (1977), Hunt and Lansman (1986) and Kosslyn (1980). All are (nontrivial) amplifications upon the production-system notion for information processing models developed by Newell and Simon (1972) (see also Newell, 1973, and Hunt and Poltrock, 1974). Our discussion will be general enough so that our remarks would apply to any of these models. For brevity, therefore, we shall simply refer to *production-system models* without further citation.

Production-system models assume two separate memory systems in the mind. These are shown schematically in Figure 1.

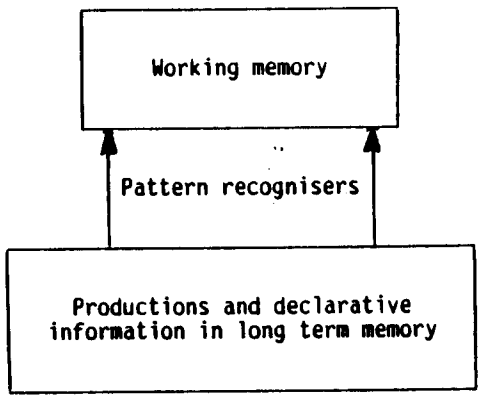

Figure 1. A production-system model of memory

Working memory is of limited space and contains information structures that are immediately at the focus of apprehension. Long term memory is a virtually unlimited bank that contains two types of information: *declarative information* about the relationships between events and concepts; and *productions* that guide action. Productions are written as pattern-action pairs, i.e., in a sort of *if-then* notation. To illustrate, a fragmentary set of rules for driving might contain the productions:

If a red light is observed, *then* apply brakes.

If a yellow light is observed, *then* examine side streets.

"If a yellow light is observed" in this example means "If a representation of a yellow light is placed into working memory." Productions, then, describe a person's *procedural* memory, what the person knows how to do. Production execution is strongly parallel. It is assumed that all productions are continually "looking at" the data structures in working memory, and that a production's action is taken when its pattern side appears in these structures. Various mechanisms have been proposed for resolving conflicts when the data in working memory matches more than one production. Again, this is a detail that need not concern us. (Further discussions and examples are provided by Hunt and Lansman, [1986] and McDermott and Forgy [1978]).

Declarative information is best thought of as information about static information expressing real world information. To continue the driving example, the information that "red lights" are "traffic signals" and are "government property" would be held in long term memory as declarative information.

What does it mean to comprehend something in this framework? Comprehension is the construction of a data structure in working memory that meets some criterion for coherence. We will be vague about what the criteria might be, but will try to illustrate by example. Suppose one hears the phrase, "The cat caught the mouse." Productions for parsing sentences and retrieving meaning would construct a data structure that would be in some sense analogous to a parsing tree. That is, we assume working memory would contain something equivalent to the propositional statement (catch [past] [cat = actor] [mouse = object]).

Our understanding of the statement would go well beyond the propositional structure, because the terms in the proposition would refer to objects richly embedded in a semantic structure. We know that cats are carnivores, that mice are animals much smaller than cats, etc. Thus, most of us could give at least a reasonable answer to the question, "Was the cat hungry?" and could certainly answer the question, "Was the cat awake?" The information required to answer these questions is implied by the original sentence, but is not contained in it. A Martian who knew only the dictionary definitions would know only that, "The cat, a middle sized carnivore that feeds on small rodents, caught the mouse, a small rodent." The Martian could deduce the implied meanings, by a sequence of substitutions of further definitions, but at what cost? The most obvious is that the Martian will have utilised working memory space to hold information a real person would hold in the much cheaper long term memory area. A slightly less obvious point is that because the information is, by definition, new to the Martian, the Martian long term memory will not contain productions that are

triggered by this data structure. A person familiar with cats and mice (perhaps a mouse lover) will have procedural knowledge that something must be done to avoid damage. Further, these procedures will be triggered immediately by the information, whereas the Martian might have to come to the same reasoning by a slower, working-memory intensive process of deduction, at greater cost to both Martian and mouse.

This is the crux of the matter. Understanding is achieved by establishing relations between objects. The relationships may be established either explicitly by constructing data structures in working memory, or implicitly by building data structures whose elements are already richly connected to other elements in long term memory. Consider an analogy to building. Presumably any frame house could be constructed from boards and nails. Prefabricated parts can greatly reduce the work involved, but if one relies on prefabricated parts, then only some buildings are possible.

What has this to do with language and thought? A language provides "prefabricated thoughts" that can be used to build a data structure for comprehension. We will refer to these as *concepts*. People try to understand a situation (build a data structure representing it) by using the concepts they already have. This is an excellent strategy because the labels for the concepts can be used within working memory to refer to very large data structures in long term memory. But sometimes the concepts cannot be formed into a structure that represents the current situation adequately. In theory, when this occurs a person should be able to fall back on a few universal primitive notions, and build a working memory structure from these universals. In practice, though, the comprehender who does not have the right labels and concepts is in as difficult a position as a building contractor who has only boards, nails, a hammer, and a saw, but no blueprint.

We shall amplify our analogy by considering different situations in which language seems to control thought. Two themes will run through our discussion. Labels (usually morphemes) categorise the world into situations where the label applies and situations where it does not. Different languages use different categorisations. How do the different categorisations influence thought? Thoughts themselves are seldom expressed by a label, they are expressed in symbolic structures; we think in sentences and paragraphs, not words. Languages differ in the rules they use to form these structures. How do these differences influence thought?

The Mechanisms for Linguistic Effects

Words. We will now amplify our use of the term "concept", which is itself one of the more vaguely defined terms in our language. (Consider the

difference between a mathematician speaking of the concept of real numbers and the advertising executive who wants a high concept campaign for a new product.)

In experimental psychology "concept" has traditionally been used to refer to the name of a set of objects (Hunt, 1962). This is too restrictive. Following Miller and Johnson-Laird (1976), Murphy and Medin (1985), and Sperber and Wilson (1986), we will stress three different aspects of a concept.(i) The first is the substitutive definition; a description of the concept, in more primitive terms, that can be substituted for the concept label in any symbol structure. For example, "small domestic feline" can be substituted for "cat" in any proposition containing "cat".

The second aspect of a concept is its relational definition. Any concept enters into relations with other concepts. To us, a "cat" is defined by its physical attributes and partly by its relation to mice. Cats are also defined by their relation to, say, the heroines of Victorian novels. The two relations depend upon different parts of the substitutive definition: the mouse relation depends upon cats as felines; the Victorian relation depends upon cats as domestic pets. Since there are objects that possess the parts of the substitutive definition to varying degrees, an individual example of a concept may be able to enter into only some of the relations that the concept normally involves. A declawed, defanged cat may be an excellent cat in a romantic novel, but a laughable cat to a mouse. Conversely, there are some unkempt, ferocious alley cats. The point is that concepts exist to be used, and when they are used, only certain of their normally defining relationships are appropriate. Any object that can play a role of a "cat" in a certain situation *is* a cat *in that context*.

The idea is that a relational aspect to concepts may strike speakers of English as unusual. We think that this is the point that Whorf was trying to make. It is probably true that concepts in every language have a definitional and a relational aspect, but languages may differ in the emphasis that they place on each aspect. Whorf claimed that the SAE languages stressed things in and of themselves, i.e., the definitional aspect. Hopi stressed the relational aspect.

Most of the terms in both the definitional and relational aspects of a concept will be other concepts. At some point though, there has to be a set of elemental, nonlinguistic terms. Presumably, the definitional terms are general across cultures, e.g., perceptions of colour. We join Schank (1972) and many others in suspecting that there are a relatively small number of relational primitives, such as "contacts", "is part of", and "strikes", etc. Surely every human group has a concept of causation, obedience and threat. What languages do is to provide elaborations of the primitives, in different, culturally-specific ways. Consider, for example, the elaboration from "strikes" to "harms" to "libels".

Words (morphemes) serve two purposes. In communication, a word is a unit that lets one person call another's attention to a concept occurring in a specific context. We are more interested in what the presence of a word in a language indicates about the lexicon of the speaker's internal thoughts. The existence of a word indicates that the speaker has an internal label for a particular concept.(ii) According to the production-system model of cognition, thoughts themselves are structures built from these labels. The working memory structures that constitute newly formed thoughts contain labels that serve as pointers to previously formed thoughts. If working memory were infinitely expandable, such a system of pointers to old ideas would be of no value, because the thinker might as well bring the old structures themselves into working memory. But working memory is limited, and so the labels are useful.

Anyone who has tried to teach statistics to undergraduates will be familiar with what we mean. The instructor comes from a culture in which terms like ANOVA are primitive labels. Most undergraduates do not, so they must drag an unwieldy collection of primitive terms into memory. More than a few of them become overwhelmed. Eventually though, they acquire the labels, become instructors, and go on to mystify subsequent generations.

The ANOVA example illustrates the confusion that can be caused when a person does not have a label for a data structure. The label is of little use, unless the person has a rule stating when the label's use is appropriate. We will call this rule the *identification function* of a concept. It is important to realise that the identification function is distinct from either the definitional or relational aspect of a concept.

An example from the Indian caste system will serve well here. In some regions of India, a person's family name indicates caste. Thus, an individual's caste can be identified as a Brahmin or a Sudra by the structure of the name. Under Indian law a person can adopt any name one wants, but no one would *become* a Brahmin by adopting a Brahmin name. In fact, there was a historic attempt to alter the relational aspect of being a Brahmin by defeating the identification function. About two hundred years ago some progressive Brahmins dropped their last names in lieu of an initial signifying the last name, so that they would not receive the special privileges that tradition assigned to them. Because the progressive Brahmins wanted a new relational aspect to the concept "Brahmin", they had to provide a new identification function. Unfortunately, conservative Brahmins also adopted the new naming convention, so the scheme was defeated, but the point remains. In fact, the actions of the conservative Brahmins illustrate the other point we wish to make. Every concept must have a unique identification function, otherwise it cannot be used.

The historic Brahmins were certainly not the only people who have confused identification functions and relational aspects of concepts. We

suggested that such confusions are particularly likely to occur in cross-cultural settings, when one culture is trying to acquire information from another. Let us call these two cultures the "observing" and "demonstrating" culture. What members of the observing culture can see directly are the situations that fit the identification function of the demonstrating culture. The conceptual reasoning of the demonstrating culture is not so obvious, and often can only be explained in terms that are themselves specific to that culture. Furthermore, the observing culture will be biased toward assimilating the situations that fit the identifying function of the demonstrating culture into their own established concepts. From the viewpoint of a designer of production systems, this is reasonable. Only trouble can result from the possession of two concepts with almost identical functions, for they will continually interfere with each other in the recognition process. Misunderstandings arise when the assimilation produces a concept in the observer that is not quite what the demonstrator intended.

We offer the following historical example from Claibourne (1983). In 587 A.D., the missionary Augustine brought Christianity to the Angles and the Saxons. He was able to explain what he meant by *Deus (Saxon-God)* and *paradise (hefen)*. The English even knew about *synne* and *hel*. However, the idea of *sanctus spiritus* was more ethereal than the pragmatic English could handle. The best that could be done was *halig gast*, which a twentieth century daughter of the Saxons defined as "Casper with a halo."

There is a serious undercurrent to this example. Apparently the hardest thing for Augustine to translate was the least perceptually vivid concept of the Trinity. "The Father" and "the Son" can be defined by universal human social relations. "Spirit" is a concept that is meaningful only to those who have already developed a supporting complex of beliefs. We shall return to this point below, but first we must consider some more points about language and thought above the level of the word.

Schema, Language and Thought

Concepts are static structures in long term memory. Thoughts are assemblies of concepts that are related to each other. Every new thought places old concepts in a new relation. Saying "Ronald owes Margaret for Libya" tells us something that may have been reasonable, given what we already knew about Margaret, Ronald and Libya, but was not dictated by that knowledge. Technically, we will speak of thoughts as data structures. These can be thought of as labeled, directed graphs in which previously learned concepts are associated with the arcs and with nodes that do not have arcs emanating from them. New thoughts, that bring old concepts into an original relation, are represented by the higher order nodes. This is illustrated in Figure 2, a graphic depiction of the "Ronald and Margaret"

example. It is often possible to present data structures in a more concise propositional notation, e.g.,

(owes [Ronald, agent] [Margaret, patient], [cause, Libya]).

We shall use either notation, whichever is more convenient at the time.

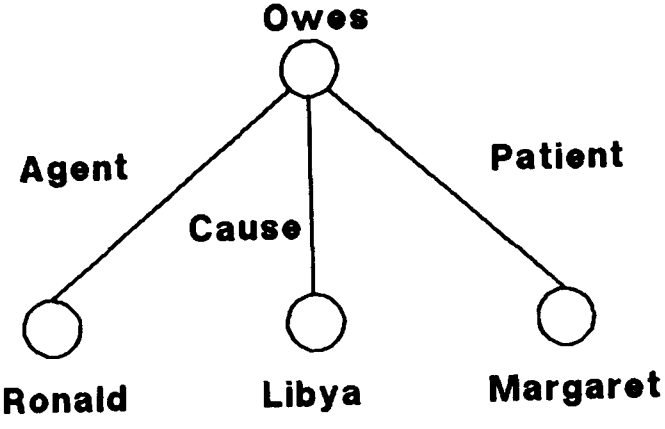

Figure 2. Graphic depiction of "Ronald owes Margaret for Libya"

In the previous section we argued that language provides the concepts used in the data structure of thought. In this section we explore two ways in which language guides the construction process. Different languages provide different devices for ordering constructions in general, rather in the way that different carpenters might use different ways to lay out their tools on a workbench. This is a rather subtle effect, so we postpone discussing it until we have examined a more striking influence, the role of schema.

Continuing the analogy to carpentry, carpenters work from a higher order plan, that directs their actions to first one part of the thing they are building, and then another. Virtually all cognitive science treatments of thought emphasise the importance of higher order units, variously called schema (the term we shall use), macropropositions, plans or memory organising procedures. These are all plans that impose order onto an imprecise or incomplete stimulus situation. Consider what higher order knowledge is required to understand the following passage.

> Lucrative offers have poured in from movie producers and tabloids that want to re-create the story of the disastrous expedition on Mount Hood, but the school that sponsored the climb is rejecting the idea as abhorrent and repulsive.
>
> Oregon Episcopal School said in a statement that it will not participate in what it termed commercial exploitation of the disaster.
>
> Seattle Times, May 25, 1986

Most people familiar with modern American journalism will have little trouble understanding the gist of the story, even if they do not know what the Mount Hood tragedy was. However, the passage would be nearly incomprehensible to anyone who did not have schema for dealing with American sensationalist journalism and the attitudes of many about their practices.

Schemas are essentially relational formulae, i.e., they state that entity x stands in relation R to entity y.[3] The terms, R, x and y can be presented at varying levels of specificity. Returning to the Mount Hood example, the schema for action and inducement dealt with unspecified persons and actions, while the schema concerning sensationalistic journalism referred to certain types of people and more precisely stated actions.

Why do we have schema? The answer "In order to achieve understanding" is not adequate, because this requires a definition of "understanding". Following the suggestions of numerous authors, we argue that schema are used primarily for two reasons: to achieve prediction and to assign causation. Since the prediction case is easiest to see, we will deal with it first.

Schema as predictive devices. One of the benefits of thinking is that manipulating a mental representation lets one avoid the hazards of manipulating the real world. For this to be successful, the thinker has to be able to construct a mental representation that accurately portrays crucial relations in the thinker's physical world. Schema are culturally satisfactory if they succeed in explaining and predicting the problems that a culture faces. Schema that fail to do so are dropped out, while schema that solve unfaced problems never occur.

To illustrate this point, we consider the linguistic development in a society of half-naked, semi-literate inhabitants of the warmer beaches on the Pacific coast of North America. Surfers speak of waves as being "hollow" or "walled". A hollow wave is one that breaks sequentially along its crest, so that the wave break may roll roughly parallel to the beach for perhaps a mile. A good surfer will ride a hollow wave just in front of the break, moving almost perpendicular to the wave's path towards the beach. By contrast, a walled wave has a nearly vertical rise, and breaks simultaneously at all

points. A wall can only be ridden directly towards the beach. These concepts have functional distinctions. Surfers can perform acrobatics on their boards while riding hollow waves, so beaches with hollow waves are considered more desirable for surfing. The ability to manipulate hollow waves, however, depends upon the design of one's surfboard. In the 1950's, before surfing technology developed, surfers did not speak of hollow and walled waves, for all waves were ridden directly toward the beach.[4]

The surfer example is an example of a situation in which a single referent can be used to describe a whole sequence of events. A surfer's statement "I rode hollow waves all day" implies a whole style of surfing in addition to specifying a wave form. The concept has obvious predictive utility; saying the waves are hollow informs the surfer of the sort of day, type, and probably intensity of surfing. Indeed, one of the benefits of having a single word for a schema is that two surfers can, briefly and succinctly, explain to each other why they are not going to work or class: "It's hollow."

Our example was intentionally graphic. However, schema may be used to order much more abstract events. In fact, one of the functions of a schema is to provide ordering for classes of situations. We have all had the experience of coming into the middle of an American "cops and robbers" movie and being able to pick up the plot almost without effort. This is because such stories are schematised. They feature a young hero who defies regulations in order to solve crimes. The hero is always defeated in the next to last reel, makes an inspired deduction, and triumphs in the last.

Detective stories with a different schema were popular in China during the 10th century Sung Dynasty. The hero was always a middle-aged magistrate who proceeded strictly according to rules, examining the crime, consulting the spirits of his ancestors, and then had the guards beat a confession out of the guilty party.

We doubt that anyone would deny that schema are used, or that different cultures use different schema. Our point is that schema have to be used, because their predictive power allows human thinkers to fix their limited computing capacity on the important parts of the situation.

Schemas as explanations of causality. It is easy to see why we need schema for prediction. Why do we need schema for causality? We will not attempt to answer this question; we simply observe that humans do not seem to be satisfied with their understanding of a situation unless they can assign causality. We shall assume that there exists a primitive (and universal) relation *cause (x,y)* which, when it can be instantiated, creates the subjective state of believing that the relation between x and y is understood. The normal way that understanding is reached is by fitting a situation to a (previously held) schema that either contains the primitive *cause* or some instantiation of it. Although the drive to find a causal relation

may be universal, what counts as a causal explanation is at least partially cultural.

Schema intended to provide causal explanations are much less constrained by the physical world than are predictive schema. Most events permit multiple explanations. Therefore, the culture has greater latitude to invent explanations than it does to invent predictions. In its time, until some very sophisticated observations were made, the concept of *phlogiston* served quite well to order the facts about combustion. Cultural freedom is even greater if the purpose of the schema is to bring either causal or predictive order to social, psychological, and in the extreme, religious and metaphysical phenomena, because in these matters the objective facts are less constraining.

How do people decide what causal schema to apply to ambiguous situations? This is an important question to us, because a production-system model has to find some cue to activate the schema that are going to be used. Evidently, at least some of the cues for activating causal schema are contained in the language. Au (1986) has reported an interesting case, the assignment of causality after hearing fragmentary sentences involving verbs of experience, such as *scare*, *upset* or *surprise*. Consider the sentence "Mohamar infuriated Ronald." Does this imply that Mohamar did something, or that Ronald is a person who is easily infuriated? (Objectively, we would be sympathetic to either explanation.) Using less political examples, Au (1986) showed that English speakers assigned causality to the agent (in our example that Mohamar did something). Au, citing her own data and related work by Brown and Fish (1983) dealing with Japanese and Chinese, has suggested that this is a cultural universal; causality is always assigned to the agent rather than the patient of an experiental verb. In another part of her study, Au showed that action verbs are more flexible. Nineteen out of twenty English speakers saw the agent as the cause of an event in *apologise* (as in "Margaret apologised to Ronald"), while none saw the agent as the cause of *congratulate*. Other action words (e.g., criticise) were seen as ambiguous. We suggest that it would be interesting to study these effects systematically, as a function of the background of the speakers. The ambiguous words are particularly interesting. We would like to know what sort of people see the agent as causing a criticism, or the patient as drawing one.

Schema that guide social relations are particularly interesting. Modern studies of communication stress the importance of a "model of the other" in social interactions. If a person x wishes person y to do action z, person x must provide y with some information that, *added to the information schema, and deductive processes y already has*, will lead y to deduce that z is an appropriate action (Sperber and Wilson, 1986). Such reasoning can lead to a very complex sequence of actions. This is

illustrated by the following account, which describes the somewhat incongruous results of combining the Western concept of banking institutions with non-Western concepts of personal obligation.

In Bombay in the early 1980's, the Maharashtra State Cooperative Bank was having difficulty collecting overdue loans from farmers. A banker's usual recourse is to the courts. The Bombay bankers adopted another strategy. Several of the managers each "adopted" an individual farmer and his loan. The adopting manager then proceeded to go on a hunger strike until firm assurance was given that the loan would be repaid. The symbolism of this act was made even more poignant by the fact the the level of seniority of the manager was commensurate with the amount of loan, so that the largest loan was adopted by the highest ranking manager. This strategy worked in Bombay. We are sure it would never have occurred to the managers of, say, the Bank of America. The point we wish to make, though, is that social behaviour (i.e., any behaviour that does not rely on physical force for its consequences) has its intended impact only because of a shared understanding and acceptance of the significance of the behaviour. People are social beings, who react to others' behaviour because they identify that behaviour as entry points into their own schema, and those schema tell them how they must respond.

What has this to do with language? We assume that the Bombay bankers spoke to each other as they developed their strategy. We also assume that they would never have adopted this strategy if they were dealing with, say, a Western shipping company. They had to talk differently about their debtors in order to plan responses appropriate to each case. If their language had not permitted this, planning would have been impossible.

Language as the entry point to schema. We do not take the extreme position that all thoughts and actions are dictated by pre-existing schema. People have the ability to construct original ideas. Our point, though, is that humans have a strong bias toward using schema to order their world. We would even maintain that most thoughts that are trumpeted as being original are, in fact, modifications of previously developed schema. Let us consider, more abstractly, what schema do and why the computational characteristics of the mind dictate the use of schema.

We have argued that "thinking" is a problem in symbolic computation. In general, there are two ways to determine the answer to any symbolic computation problem: by applying an algorithm that builds an appropriate symbol structure in working memory; or by looking up an answer and placing it in working memory. No general rule can be given to say that one method is better than the other; it depends upon the relative costs of computation and "lookup". This can be illustrated by the ways in which transcendental functions have been "calculated" over the years. The

common transcendental functions (sine, cosine, logarithm, etc.) can be approximated to any desired degree by computations that, although conceptually simple, are tedious for a human to perform. So, prior to about ten years ago, people looked up the values of transcendental functions in tables. Today most people who deal with transcendental functions use hand calculators and computers, recomputing the functions as desired. The relative costs of computing and "lookup" have changed.

Schemas function in a manner analogous to tables. They are devices for shifting the burden on a computation from symbol manipulation to "lookup". Tables, of course, are an extreme example, for they provide for exactly one, context free solution. (The natural logarithm of 2.0, to five decimal places, is always 0.69315.) Perhaps a better example would be a table of forms for integration. It is possible to do symbolic integration on a computer, but there is still room for a book (i.e., a set of schemas) of forms.

We doubt that anyone would seriously argue with the propositions that schema are important in human reasoning and that many schema are culture specific. But what has this to do with language? Our argument is that the symbols contained in a schema's symbol structure are the internal "mentalese" terms for a person's concepts. While we would not argue that the named concepts in a person's language and the concepts of thought are exactly coterminous, we do argue that for any term in the external language there must be an internal concept. This concept will appear as a primitive term in many memorised schema, and will point to these schema when it (the concept) appears in a working memory structure. Those schemas that are most activated by current contents of working memory will be the schema used to interpret those contents. The point is simply that the initial stages of any pattern recognition system must be "bottom up", starting with the language elements themselves.

This can be shown in an elegant manner by considering situations in which the linguistic cues themselves can only be interpreted by the use of schema. Clark and Clark (1979) have pointed out that American English is rife with "verbified" nouns, such as "Rover treed the postman." The Clarks argued that a noun can be verbified only if the nouns named point to an unambiguous schema that contains a relation not named in the utterance. For instance, what relation could possibly exist between Rover, a tree, and a postman? This facility in English can be used to invent instant, highly culture-specific schema. We offer two further examples, to show how the languages and schema of a subculture determine the invention of a new term, which can then be used to construct still further new schema.

In American research universities some professors are peripatetic. One of our colleagues said "They are training me to Boston." Because of the schema associated with this particular speaker, we knew at once that (name withheld) was being transported by rail. The example is a strong

case of the use of schema, since "training" is itself a verb in a different context. Most of our colleagues will have no trouble understanding this illustration. But what about "Congress had Christmas-treed this bill," a phrase used by the leader of the Potomac tribe? Can speakers of Academic English understand this? Only if they have pre-existing schema of a piece of legislation as a gift for everyone.

The last example is, in fact, a serious one. A number of years ago Elliot Richardson, then Secretary of Defense, remarked that until he came to the Pentagon he had not heard "Christmas tree" used as a transitive verb. Since that time, though, we have observed several cases of its use, and of its amplification, both in the press and in conversation with Washingtonians. It seems an interesting example of how linguistic terms are used to develop and maintain a concept.

Language and the construction of thought. Our last illustration was an example of how data structure, that was invented to describe a particular situation, proved useful enough to graduate to the status of a schema in long term memory. Most of our working memory data structures are transient. The language we speak may still aid in their construction, by facilitating the way in which we keep track of the concepts we are trying to fit together. It is important to realize that this is a relativistic statement; we do not believe that there are thoughts that are completely restricted to any one language. We do believe that the mechanics of the mind interact with the characteristics of a language to make certain structures preferable in one language, and other structures preferable in another.

We shall offer some examples of what we mean. However, we have found it much more difficult to do this than to construct examples of schema or concept use, because the relevant data are simply not present. There is a theoretically justified reason for this. We want to discuss how language influences the mechanics of thought, not the contents. By definition, the mechanical aspects of thinking are not available to conscious experience, whereas the contents are. Since schema contain content, we can observe them simply by knowing (or being told about) their existence. On the other hand, observing the mechanics of, say, memory scanning, requires a sophisticated experimental situation. By and large, such observations have not been taken except within the context of the English language. Perhaps this paper will inspire the necessary cross-cultural experimental psychology.

One of the most important mechanisms used to tie discourses together is coreference. Consider the statement,

> The Boyars hated Ivan because he had abrogated their ancient rights and privileges.

The word "he" appears as the agent in a proposition ("he had abrogated their ancient rights") that is subordinate to the main proposition, that "the Boyars hated Ivan". In order to understand the sentence a comprehender must know that "he" refers to Ivan. This is called an anaphoric reference. Resolving the reference requires a search of working memory for a possible referent at the time that "he" is encountered.

Languages differ in the amount of support provided for anaphoric reference. One of the most widespread examples is the presence of the "tu" and "vu" forms in most languages (informal and formal ways of saying "you"), but not in modern English. How should the following discourse be understood?

> When the woman answered the doorbell, she found her son there, accompanied by a policeman. She immediately said 'Will you please tell me what is going on here?'

Who is the woman speaking to? There is no way to know, in English, because the pronoun "you" does not indicate status. In Spanish (and many other languages) the mother would use the "tu" form of the second person pronoun to speak to the child, and the "vu" form to speak to the policeman. In other cases English is the less ambiguous language. The English third person pronoun distinguishes gender: he or she. Turkish pronouns do not. Research on English (Ehrlich, 1980) has shown that speed of comprehension of anaphoric referents depends upon the ambiguity of the referring term. A straightforward extrapolation leads us to expect analogous cross-linguistic influences. It would also be interesting to investigate usage. Do different languages evolve different ways of saying the same thing, in order to minimise the burden on working memory?

A current controversy about the Whorfian hypothesis offers a further illustration of the point we are trying to make here. Bloom (1981) observed that Chinese does not contain a structure analogous to the English subjunctive. He reasoned that, therefore, Chinese should have difficulty comprehending counterfactual statements. English counterfactuals can be stated using the subjunctive, "if X were the case, then Y would follow." A Chinese speaker would have to say "X is not the case. If X, then Y." In our terms, the English statement of the counterfactual can be expressed in a single propositional structure: implies (X, false) Y. The Chinese version of the statement involves two propositions: (not [X] implies [X, Y]). Research in English has shown that the number of propositions in a statement is a powerful determinant of the comprehensibility of that statement (Kintsch and Keenan, 1973). Therefore, according to Bloom, Chinese should have difficulty with counterfactuals.

Bloom's analysis has been mired in controversy. Bloom (1981) gathered evidence that Chinese speakers in Hong Kong were more likely to misunderstand a counterfactual story than were English speakers. Au (1983) and Liu (1985) report contrary evidence. The controversy over the validity of the evidence reported by these investigators has been framed in terms of the adequacy of Bloom's translations, and the difference between concrete and abstract stories, both of which have been used at various points in these studies. We do not see the issue as whether or not Chinese can understand counterfactuals at all. (Since China has been famous for its subtle diplomacy for more than three thousand years, we are sure that Chinese can, indeed, reason counterfactually.) The question is, at what cognitive effort? What we think would be more fruitful would be an examination of the demands that counterfactual reasoning in each language makes on working memory. It is possible to measure these demands during reading and speaking (Daneman and Carpenter, 1980; Daneman and Green, 1986). Perhaps these procedures should be applied to evaluate Chinese and English counterfactual reasoning, instead of relying on the cruder measure of absolute comprehension.

The Limits on the Whorfian Hypothesis

It appears to us that very many people, especially those with bilingual experience, believe that Whorf's hypothesis is correct. On the other hand, controlled experimental and observational studies have generally failed to confirm the hypothesis. This is a serious state of affairs because, as we have argued, some form of Whorf's hypothesis is implied by today's models of the mind. If the hypothesis cannot be confirmed, then there is cause to doubt the models. It is appropriate, therefore, to consider where one would expect the Whorfian hypothesis to apply.

Our version of Whorf's hypothesis is definitely a hypothesis of linguistic relativity, rather than linguistic determinism. Concepts and concept-based schema colour thought, but they do not absolutely determine it. In particular, we would expect the application of the Whorfian hypothesis to be modulated by two factors. One is pragmatic. Concepts and schema evolve to solve mental problems. To the extent that human cultures face common problems that permit only a restricted range of solutions, they should develop nearly equivalent linguistic concepts, including their related mental schema. The most obvious schema that spring to mind are the Piagetian concepts of conservation and object constancy which, in some form or another, do appear to be widespread. (Indeed, it is hard for us to imagine normal thinking without these schema.) Similarly, the Inuit, the Australian Aborigine, and the Western European all have to cope with a world in which objects fall down in a predictable fashion. This does not

mean that everyone, in every culture, has a complete understanding of the physical world. It does mean that the laws of physics place rather strict limits on the range of acceptable mental models of the physical universe. We do not maintain that all cultures have an intuitive understanding of acceleration, velocity, friction, and inertia. We do maintain that functioning members of every culture know that there are things that can be moved, that the bigger they are, the more push it takes to move them, and that where or when a rolling stone stops depends both on what it is rolling over, and whether it is rolling uphill or downhill.

The second limitation on linguistic (and cultural) control of thought is based more on people's flexibility to deal with new problems than upon their need to cope with reality. We have argued that the linguistic primitives of a person's language provide a set of tools for thought. We believe there are other tools, based on perceptual and perhaps primitive conceptual capabilities that can be manipulated "in the mind's eye" (or ear), but that have no direct tie to language. To illustrate this, consider the two "form board" problems presented in Figure 3. In each case the question to be solved is whether two figures can be fitted together, as in a jig-saw puzzle, to form a single piece. It is easy to see that the answer is "yes" for the two pieces at the ~~top~~ bottom. Probably most people more or less said to themselves "Fit the triangle projection on the right-hand figure into the triangular indentation on the left-hand figure." The problem at the ~~bottom~~ top is a bit harder. The lines are irregular, so that very few people are likely to name parts that can be brought together at the juncture. Nevertheless, most people can solve such problems, probably by manipulating visual images without any linguistic support (Pellegrino and Kail, 1983).

The extent to which nonlinguistic images can be manipulated is apparently quite limited. There is a great deal of research showing that people cannot construct infinitely large images (Kosslyn, 1980). Furthermore, there seem to be strong individuals differences in this ability (Pellegrino and Kail, 1982). Nevertheless, nonlinguistic problem solving can be done, and the facility of doing so is certainly accessible to all cultures. Thus, any perceptual task that can be solved by a simple nonlinguistic strategy should be within the limits of any culture. This explanation can be offered for data obtained by Heider and Olivier (1972). They studied colour memory and colour sorting in the Dani, a group in New Guinea who, although they have colour vision, do not have any linguistic terms for colours other than "black" and "white". Heider and Olivier's tasks were relatively simple. We argue that these tasks could be solved by purely perceptual strategies that would be subject to any biases that are involved in human perception and/or memory for sights.

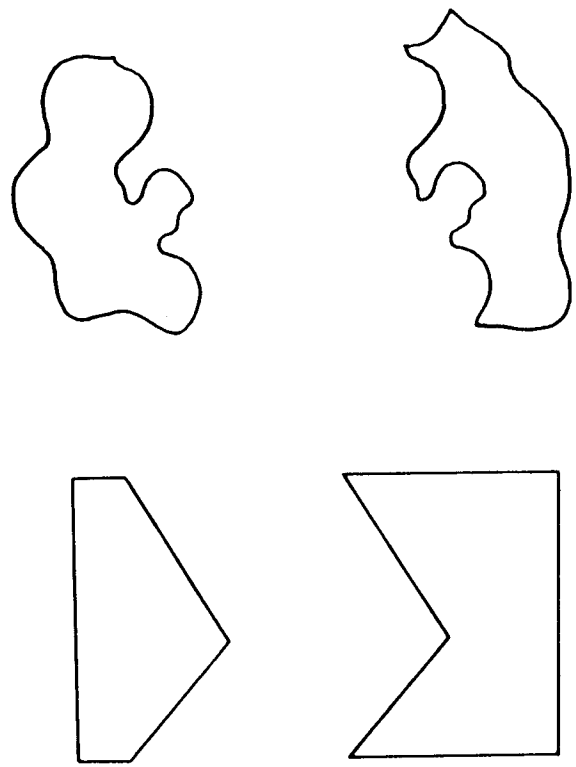

Figure 3. Two formboard problems

This is a point at which we differ from Whorf. Readers may have been struck by the parallel between the surfer example we presented earlier and Whorf's own example of the many names Inuit have for snow. Whorf argued, in his more extreme statements, that the existence of a labeled concept actually changed one's perception of an object. We do not think so. In fact, readers of this chapter are now quite capable of applying the identification function for hollow and walled waves, even though they may be unable to appreciate the esthetic and cultural importance of the distinction. (See Brown, [1958] for a similar point.)[5]

Abstract Thinking, Explanations, and Cross-Cultural Understanding

At several points we have remarked that both our schema and our language provide causal explanations for events. Why should the machinery of the mind demand causal explanations, and how does this affect the interaction between language and thought?

Causal explanations impose a hierarchical ordering on schema. This makes it possible to treat new situations as special cases of a general rule. Consider the banking example discussed earlier. The immediate problem for a banker is to collect the payments from specific individuals. Rather than ask, "How can one enforce paying of debts on X?", which might require extensive knowledge of X's individual characteristics, it is easier to solve the abstract case of enforcing social contracts in general, and then to specialise the solution. Since the linguistically described components of the abstract case (here "people", "obligations", etc.) are themselves abstractions; they can be described succinctly, without overburdening working memory with details. In order to describe the forces that act on abstractions, however, some concept of causality must be constructed. The wind pushes a sailboat, but a force accelerates an object.

The use of causal models forces a compromise between two aspects of mental computation. One is the need for accurate pattern recognition. A thinker has to be able to recognise the concepts and schema that are appropriate in a given situation. The identification functions of concepts and schema must be distinct enough so that they are not confused. If an identification function cannot distinguish between two different situations, the schema involved must be general enough so that the appropriate actions can be determined as specialisations of a general case. The argument for generality is straightforward; a small number of consistent schema can be used to deal with a wide variety of situations. The flight of footballs, basketballs, arrows, and rockets can all be treated as special cases of a "physics of acceleration" schema.

The argument against the use of abstract, general schema is that the translation process itself may overwhelm the thinker's attentional capacity. Whether or not this is true will depend very much upon how the thinker codes the problem, i.e., what language the thinker uses. If the language is efficient the process of encoding a concrete problem as an abstract one will proceed efficiently. If the language is not efficient, the thinker's computational capacities may be overwhelmed. Whether or not the thinker would think in terms of a single, highly abstract schema or in terms of more limited, possibly confusable ones depends, as so many other things do in thought, upon the trade-off between computation and pattern recognition.

Cultures develop efficient languages for coding and treating problems that they encounter frequently. This allows them to do two things: treat a

specific problem efficiently and give a causal explanation for classes of frequently occurring problems. At the same time, though, the development of a language for talking about a problem may interfere with the ability to see a problem in a new light. The interference can be of two sorts. Reformulating the problem may require a great deal of computation, so that the problem solver is unable to go from the concrete case to the causal schema because one's language is not efficient enough. Whorf's remarks about the Hopi conception of time exemplify this. When Whorf said that he could not envisage Hopi thinking about the physics of time, he surely did not mean that he was utterly incapable of doing so. English probably does have enough words to construct propositions expressing the Hopi concepts. The problem is that the resulting symbol structures would be so unwieldy that no human could understand them. To think about the Hopi concept, the English speaker would have to invent new concepts that would be ... Hopi.

The Direction of Future Research

No one would claim that modern cognitive science had found out all there is to know about the mind. However, the contrast is not between knowing all about the mind and knowing nothing about it. There is enough truth to the production-system model of mental action to use it seriously in considering some of the classic questions about language and thought. We, like Whorf, believe that the language a person speaks strongly influences the cognitive structures they can build.

Our arguments presented here have been buttressed by principled examples, taken largely from historical accounts, anthropological field work, and even the daily newspapers. Such a line of argument is somewhat surprising coming from experimental and social psychologists, since so much of our own discipline is oriented toward gaining knowledge from controlled experimental studies. We would certainly welcome the use of experimental paradigms to test our ideas, but we are not all all sure that they are really feasible. Language is only going to control thought if the thinker really knows the language.

This means that the appropriate subject is a person who is immersed in a particular culture. Teaching someone an artificial language, or teaching them a natural language that they are not required to use heavily, simply does not implant the language well enough to make it possible to test any of our ideas about the language-thought interaction. A possible exception to this would be to study programmers who work in markedly different languages, such as LISP contrasted to PASCAL. We know of few such studies, and those that have been reported seem to be oriented toward fairly applied questions in the psychology of programming rather than general

questions about language and thought. With some regret, we conclude that the appropriate way to investigate questions about language and thought is going to be to contrast language and problem solving in different cultures. Enforcing experimental control will all too often throw the baby out with the bath water.

There are essentially two ways that such contrasts could proceed. One is to choose theoretically motivated problems, and present them to two different language groups. Assuming that the groups are otherwise equated, one could then apply the usual logic of an experimental paradigm to the study of language and thought. We have already suggested that the controversy over Chinese and English counterfactuals could be viewed in this light, providing that one concentrated on the information processing aspects of comprehension, rather than on the level of comprehension.

To amplify this point, information processing studies of psycholinguistics have very largely been studies of the psycholinguistics of English. The Whorfian hypothesis can be thought of as a statement that information processing is qualitatively different in different languages. The variations and invariances in information processing across languages should be predictable from a model of the mechanics of thought.

While we do see a role for experimental studies, we do not think that they will ever replace the insightful analyses of principled examples. Here we go clearly beyond psychology (narrowly conceived) into comparative linguistics, cross-cultural psychology and anthropology. Historical analysis of the development of concepts and languages containing them could prove to be an important source of evidence for or against our ideas. The necessary data gathering, however, would have to be based on professional historical analyses, not a collection of anecdotes that two psychologists happen to know something about. We will regard this paper as more than a success if it fosters replies that either show that the modified Whorfian hypothesis is true, or forces us to rethink our ideas about the computations of the mind.

Notes

1 We shall not attempt to develop a complete theory of concepts here. The approach we have taken is closely related to that of Miller and Johnson-Laird (1976). Our account of how concepts are combined to form thoughts has borrowed heavily from Schank's (1972) notions of conceptual dependency analysis.

2 The converse is not true. A person might have a well formed concept with an internal label, but no word for it. We suspect that this situation is rare.

3 The argument could be put more formally, by asserting that schema were always expressed in a propositional form. In such a case schema could be placed in a lattice of abstractions, by observing progressive restrictions on free variables. Presumably certain primitive terms exist for defining elemental schema. These constitute the internal language of thought and are presumably culture-universal. Various attempts have been made to deduce this internal language. See Schank (1972) for one such attempt. We shall not attempt such a construction here.

4 We thank a veteran surfer, Mr. Robert Hunt of San Clemente, California, for this example.

5 This is not the say that perception is unaffected by schema. Observers, dating at least from William James, have noted that the physical world is too varied to be attended to all at once. Schema certainly do influence perception by specifying certain things to attend to. To offer another cultural example, at political rallies the politicians and the security guards will apply different schema for crowd viewing, and no doubt will see different things. But this does not mean that the one could not adopt the other's point of view, just that they probably will not.

REFERENCES

Anderson, J.R. (1983). The architecture of cognition. Cambridge, MA.: Harvard University Press.

Au, T.K-F. (1983). Chinese and English counterfactuals: The Sapir-Whorf hypothesis revisited. Cognition, 15, 155-187.

Au, T.K-F (1986). A verb is worth a thousand words: The causes of and consequences of interpersonal events implicit in language. Journal of Memory and Language, 25, 104-122.

Bloom, A.H. (1981). The linguistic shaping of thought: A study of the impact of language on thinking in China and the West. Hillsdale, N.J.: Lawrence Earlbaum.

Brown, R. (1958). Words and things. Glencoe, IL.: Free Press.

Brown, R. & Fish, D. (1983). The psychological causality implicit in language. Cognition, 14, 237-273.

Clairbourne, R. (1983). Our marvelous native tongue: The life and times of the English language. New York: Times Books.

Clark, H.H. & Clark, E. (1979). When nouns surface as verbs. Language, 55, 767-811.

Daneman, M. & Carpenter, P.A. (1980). Individual differences in working memory and reading. Journal of Verbal Learning and Verbal Behaviour, 19, 450-466.

Daneman, M. & Green, I. (1986). Individual differences in comprehending and producing words in context. Journal of Memory and Language, 25, 1-18.

Ehrlich, K. (1980). Comprehension of pronouns. Quarterly Journal of Experimental Psychology, 32, 247-255.

Hayes-Roth, B. & Hayes-Roth, F. (1977). A cognitive model of planning. Cognitive Science, 19, 275-310.

Heider, E.R. & Olivier, P.C. (1972). The structure of the colour space in naming and memory for two languages. Cognitive Psychology, 3, (2), 337-354.

Hunt, E. (1962). Concept learning: An information processing problem. New York: Wiley.

Hunt, E. (1983). On the nature of intelligence. Science, 219, 141-146.

Hunt E. & Lansman, M. (1986). A united model of attention and problem-solving. Psychological Review, 93, 446-461.

Hunt, E. & Poltrock, S. (1974). The mechanics of thought. In B. Kantowitz (Ed.), Human information processing: Tutorials in performance and cognition. Potomac, Maryland: Lawrence Erlbaum.

Kintsch, W. & Keenan, W. (1973). Reading rate and retention as a function of the number of propositions in the base structure of sentences. Cognitive Psychology, 5, 257-274.

Kosslyn, S.L. (1980). Image and mind. Cambridge, MA.: Harvard University Press.

Liu, L.G. (1985). Reasoning and counterfactually in Chinese: Are there any obstacles? Cognition, 21, 239-270.

McDermott, J. & Forgy, C. (1978). Production system conflict resolution strategies. In D.A. Waterman & F. Hayes-Roth (Eds.), Pattern directed inference systems. New York: Academic Press.

Miller, G. & Johnson-Laird, P.N. (1976). Language and perception. Cambridge: Cambridge University Press.

Murphy, G.L. & Medin, D.L. (1985). The rule of theories in conceptual coherence. Psychological Review, 92, 289-318.

Newell, A. (1973). Production systems: Models of control structures. In W. Chase (Ed.), Visual information processing. New York: Academic Press.

Newell, A. (1980). Physical symbol systems. Cognitive Science, 4, 135-183.

Newell, A. & Simon, H.A. (1972). Human problem solving. Englewood, N.J.: Prentice Hall.

Pelligrino, L.W. & Kail, R. Jr. (1980). Process analysis of spatial aptitude. In R. J. Sternberg (Ed.), Advances in the psychology of human intelligence, Vol. 1. Hillsdale, N.J.: Lawrence Earlbaum.

Pylyshyn, Z.W. (1984). Computation and cognition. Cambridge, MA.: Bradford Books, MIT Press.
Schank, R. (1972). Conceptual dependency: A theory of natural language understanding. Cognitive Psychology, 3, 552-631.
Sapir, E. (1941). Language, culture, and personality: Essays in memory of Edward Sapir, (edited by L. Sapir). Menasha, WI.: Sapir Memorial Publication Fund.
Sperber, P. & Wilson, D. (1986). Relevance: Communication and cognition. Oxford: Basic Blackwell.
Tuchman, B.W. (1978). A distant mirror: The calamitous 14th century. New York: Ballantine.
Whorf, B.L. (1956). Language, thought and reality. Cambridge, MA.: MIT Press.

ALPHABETIC LITERACY AND BRAIN PROCESSES

Derrick de Kerckhove

The object of this chapter is to raise a basic question concerning the underpinnings of Western culture. Did the fully phonetic alphabet invented by the Greeks circa 750 B.C. (for a discussion of a possibly earlier date, see Naveh, 1982) and still used today in Greece (and in the rest of the West in its Latinised version), have a conditioning impact on the biases of specialised brain processes in our culture? Could the alphabet have acted on our brain as a powerful computer language, determining or emphasising the selection of some of our perceptual and cognitive processes? This question has already been raised in terms of hemispheric specialisation by Joseph Bogen (1975, p. 29):

> It is likely that some anatomical asymmetry underlies the potential for hemisphere specialisation; but it is also clear that the extent to which capacities are developed is dependent upon environmental exposure. Although humans of any culture, so far as we know, have the potential for reading and writing, many remain nonliterate and thus fall short of acquiring the most special of left-hemisphere functions. Conversely, we can readily comprehend the concept of a society in which 'right-hemisphere illiteracy' is the rule. Indeed, our own society (admittedly complex) seems to be, in some respects, a good example: a scholasticised, post-Gutenberg, industrialised, computer-happy exaggeration of the Graeco-Roman penchant for propositionising.

During the seventies, split-brain research gave rise to a number of scientific and popular theories, among which a book by J. Jaynes (1976) called "The Origin of Consciousness in the Breakdown of the Bi-Cameral Mind" In it, Jaynes claimed that different aspects, including the invention and development of writing, of ancient Greek culture were responsible for the rise of self-consciousness in the Western world. Although this notion had already been entertained by different avenues of classical, psychological and anthropological scholarship, especially by Innis (1950),